

**EXHIBIT 29
[FILED UNDER SEAL]**

HIGHLY CONFIDENTIAL

**UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
SHERMAN DIVISION**

The State of Texas, et al.,

Plaintiffs;

v.

Google LLC,

Defendant.

Case No. 4:20-cv-00957

Hon. Sean D. Jordan

Special Master: David T. Moran

EXPERT REPORT OF DR. JOHN CHANDLER, Ph.D.

JUNE 7, 2024


DR. JOHN CHANDLER, Ph.D.

HIGHLY CONFIDENTIAL

Table of Contents

	Page
I. Background, Qualifications, and Experience.....	1
II. Assignment.....	4
III. Summary of Opinions	6
IV. Origins and Overview of Open Web Display Advertising.....	9
A. The Advent of Digital Marketing	9
B. Research Sources.....	11
C. Key Definitions	12
D. The Marketing Funnel	17
V. Website Display Advertising and Other Digital Marketing Channels	19
A. Opinions 1–2	19
B. Display Advertising	19
C. Other Digital Marketing Channels.....	21
1. Search Advertising	21
2. Social Media	22
3. Digital Video.....	25
4. In-App Advertising	26
5. Email	27
D. Differentiating Website Display from Other Channels Based on Audience, Activity, and Medium	27
E. Differentiating Website Display from Other Channels Based on Price	29
VI. The Rise of Programmatic Ad Auctions and Placements.....	31
A. Opinions 3–5	31
B. Introduction	32
C. Differentiating Programmatic Auctions from Other Forms of Ad Sales	34
D. The Advantages and Basic Rules of Programmatic Auctions	35
VII. The Ad Tech Ecosystem.....	40
A. Opinion 6	40
B. Introduction	41

HIGHLY CONFIDENTIAL

C.	Publisher Ad Servers	44
D.	Advertiser Ad Buying Tools.....	49
1.	Overview	49
2.	The Bifurcation of Tools for Large Advertisers and Small Advertisers.....	50
E.	Ad Exchanges	53
VIII.	Google's Position in the Digital Marketing Space.....	54
A.	Opinions 7–10.....	54
B.	Introduction	55
C.	A Brief History of Google's Acquisitions and Rise to Power in the Ad Tech Industry.....	56
1.	Google's Ad Tech Tools and Platforms.....	65
2.	DFP (GAM)—Publisher Ad Server.....	67
3.	AdX (GAM)—Exchange.....	67
4.	Google Ads—Ad Buying Tool for Small Advertisers	68
5.	DV360—Ad Buying Tool for Large Advertisers.....	68
6.	Google's Other Digital Ad Tools and Platforms	68
D.	The Potential Conflicts of Interest and Self-Preferencing Opportunities Inherent in Google's Position.....	68
IX.	The Importance of Information and Transparency in the Programmatic Display	69
A.	Opinions 11–16.....	69
B.	Introduction	70
C.	Tracking Activity Online.....	75
D.	Google's Privacy Sandbox Program	78
E.	Measuring Marketing Effectiveness.....	81
F.	Summarizing the Data in the Auction Process	83
G.	The Potential Effects of Information Imbalances, Exclusions, Failure to Disclose, and Misrepresentations	86
1.	The Buy Side Wish List	86
2.	The Sell Side Wish List.....	87
3.	The Exchange Wish List.....	88

HIGHLY CONFIDENTIAL

X.	An Analysis of Google's Activities in the Website Display Ecosystem	89
A.	Opinion 17.....	89
B.	Introduction	90
C.	Google Failed to Disclose Changes to Auction Rules	90
D.	Google Failed to Provide Equal Access and Functionality to Competitors and Customers	93
E.	Google Restricted the Flow of Information.....	94
F.	Google Exploited Conflicts of Interest and Engaged in Self-dealing....	95
G.	Google Engaged in Experiments that Exploited its Users	96
XI.	Reservation of Rights.....	102

HIGHLY CONFIDENTIAL

I. Background, Qualifications, and Experience

1. I am a professor of marketing and have been a practitioner of marketing for twenty-five years. I have worked in analytics and data science since 1999 with a primary focus on digital marketing. As I will elucidate below, I have particular experience in the areas of digital marketing that I have been asked to assess. My academic work is centered on applied data science.

2. Regarding my academic credentials, my primary academic affiliation is at the University of Montana, where I am a Clinical Professor of Marketing and Ruff Family MS in Business Analytics Faculty Fellow. Additionally, I am a Visiting Professor of Marketing at Universidad ORT Uruguay and an adjunct professor at the University of San Diego. I earned a Doctorate in Statistics from the University of Montana (2010), a Master's Degree in Mathematics from the University of Washington (1999), and an Honors Bachelor's Degree in Mathematics from Middlebury College, magna cum laude (1996).

3. Upon graduation from the University of Washington, I began working in marketing analytics with Avenue A, an advertising agency. At the time, Avenue A was the largest digital marketing agency in the world. I was the primary analyst for dozens of advertising clients who spent millions of dollars on marketing. I helped develop fundamental techniques for analyzing digital marketing performance. In 2000, we formed aQuantive, a holding company with Avenue A as one division. I joined the newly created second division, Atlas DMT, a provider of digital marketing technology to the largest advertising agencies in the world. Atlas was the second-largest provider of advertiser tools, especially advertiser third-party ad serving, behind only DoubleClick, our primary competitor. DoubleClick was acquired by Google in 2008, and DoubleClick formed the backbone of Google's display advertising products and services.

4. At Atlas, where I became the sole Principal Analyst in the company's history, I had three main responsibilities: client analytics, thought leadership, and research and design on new products and product features. As a client analyst, I provided custom analytics consulting to our largest advertisers and agencies. From 2000 through 2007, I worked with dozens of agencies and hundreds of advertisers at all scales. I analyzed data from media plans across all marketing channels including display, search, video, and email. My responsibilities included helping these advertisers understand their marketing performance, optimize their spending, and determine how consumers were being influenced by their advertising. This work gave me a deep and foundational understanding of digital marketing, particularly how it was practiced on the "buy side," represented by advertisers. In particular, I worked with data from all major publishers including stand-alone websites, ad networks, and, ultimately, exchanges. I began working with data from the network that ultimately became AdX starting in 2000.

HIGHLY CONFIDENTIAL

5. My thought leadership involved writing white papers for the industry and giving talks at industry conferences. These white papers were highly scrutinized by my partners and competitors. The first white paper I wrote, “Online Holiday Shopping Patterns Revealed,” appeared on the front page of *The Wall Street Journal*¹ and was the first research to identify the phenomenon of “Cyber Monday.” Often, my team transformed thought leadership projects, via a research and design process, into products or product features for our technology platform, Atlas.

6. In 2003, my team and I launched Drive Performance Media (DrivePM), which leveraged our technology and analytics to create an ad network, arbitraging unsold digital advertising inventory. I was responsible for building the first large-scale ad-allocation engines. This piece of software was responsible for solving the problem of determining which ad from which advertiser was shown for an available impression. While working for DrivePM, I did foundational work in the data science underpinning advertising networks and exchanges. I was responsible for building models estimating the performance of inventory and forecasting the volume of inventory we would have to sell. I built large-scale non-linear optimizers to maximize performance for advertisers and publishers, subject to thousands of constraints created by the different types of advertising deals. The creation of DrivePM presaged the programmatic display revolution and allowed me to work on data science problems fundamental to both demand-side platforms (DSPs) and supply-side platforms (SSPs). I was the first analyst to work on DrivePM and I had shared responsibility for developing algorithms generating tens of millions of dollars in profit.

7. I was the lead researcher on Atlas’s tool “Engagement Mapping,” which, when it launched in 2008, was the first large-scale multitouch attribution tool in marketing technology. Multitouch attribution is a technology that allows advertisers to apportion the credit for sales across multiple touchpoints in the digital marketing funnel.² My dissertation research involved building Cox proportional-hazards models with time-varying covariates to estimate the appropriate weights for these touchpoints. Multitouch attribution ultimately became the gold standard by which digital marketing campaigns were measured.³ The research white paper I wrote, “Measuring ROI Beyond the Last Ad,”⁴ was a reference work for subsequent tool developers.

8. aQuantive was acquired by Microsoft in 2007, and I joined the research team at Microsoft Advertising, as part of Microsoft’s Advertiser and Publisher

¹ The Wall Street Journal. “Consumers Are Likely to Turn to Web For Holiday Shopping, Analysts Say” (November 21, 2001). Accessed on June 3, 2024. <https://www.wsj.com/articles/SB10006378597369697160>.

² Click Z. “How One Advertiser Uses Microsoft Engagement Mapping” (May 12, 2008). Accessed on June 3, 2024. <https://www.clickz.com/how-one-advertiser-uses-microsoft-engagement-mapping/63267/>.

³ Digiday. “WTF is multi-touch attribution?” (August 30, 2019). Accessed on June 6, 2024. <https://digiday.com/marketing/what-is-multi-touch-attribution/>.

⁴ Chandler-Pepelnjak, J. “Measuring ROI beyond the last ad.” *Atlas Institute Digital Marketing Insight*. 2009. pg. 1-6.

HIGHLY CONFIDENTIAL

Solutions (APS) division. At this time, it was part of Microsoft's strategy to provide tools for large advertisers, though it ultimately sold the assets to Facebook and stopped participating in the business of selling tools to large advertisers.⁵ As part of my work with Microsoft's APS, I was the data scientist responsible for the billions of daily impressions we released to advertising exchanges. I built models that created price floors for advertising inventory that we released into a real-time bidding auction environment managed by a third party. As part of this work, I became intimately familiar with auction dynamics, helping to write code underpinning our participation in auctions and real-time bidding (RTB) environments.

9. In 2010, I became Research Director at Microsoft TV. In this capacity, I was responsible for the analysis of a data set of four million households having a cable set-top box with our software on it. We used the information from these set-top boxes to arbitrage television advertising inventory.

10. After leaving Microsoft in 2012, I founded a data science consulting company, Data Insights, which has provided enterprise-class marketing and data science to dozens of clients. Our larger clients include [REDACTED]

[REDACTED] In addition to these clients, we have worked with many smaller clients. In the course of this consulting work, we have built statistical models and methodologies for a wide variety of business applications. Many of these projects relate to marketing analytics. I have worked with digital marketing data consistently during the entire period I have consulted.

11. I have had numerous consulting engagements that allowed me to work with ad tech companies, typically working on marketing measurement. These relationships have afforded me perspectives similar to those that I enjoyed at Atlas—working across dozens of advertisers and seeing the data they receive, the challenges they face, and understanding their position in the complicated digital advertising landscape.

12. I have remained professionally and academically engaged in the fields of advertising, marketing, marketing analytics, marketing measurement, the application of data science to marketing, ad tech, and the ad tech ecosystem. I have worked with companies on the buy side with annual marketing budgets ranging from under \$100,000 per year to those with budgets in the billions. I have worked with four of the companies whose advertising budget is in the top ten in the US as measured by Statista.⁶

⁵ Tech Crunch. "Facebook Confirms It Will Acquire Atlas Advertiser Suite from Microsoft To Close The Ad Spend Loop" (February 28, 2013). Accessed on June 3, 2024.

<https://techcrunch.com/2013/02/28/facebook-acquires-atlas/>.

⁶ Advertising Age. "Largest advertisers in the United States in 2022 (in billion U.S. dollars)." Chart. June 26, 2023. Statista. Accessed June 3, 2024. <https://www.statista.com/statistics/275446/ad-spending-of-leading-advertisers-in-the-us/>.

HIGHLY CONFIDENTIAL

13. In addition to my professional experience, I have academic experience relevant to advertising ecosystems, advertising measurement, and programmatic display advertising. In my dissertation, I applied the tools of statistical learning to web-scale data sets and built advanced attribution models on millions of records to understand marketing effectiveness.⁷

14. In 2016, I co-authored Algorithms for Data Science, one of the first books on data science as a subject and the first textbook to collect and illustrate fundamental algorithms across both machine learning and statistics.⁸

15. I have based my opinion on my education in statistics and my work experience in marketing and marketing analytics. I am being compensated for my work in this case at my customary rate of \$750 per hour. All the opinions I offer herein I hold to a reasonable degree of professional and scientific certainty. My curriculum vitae is attached as Appendix A and a full list of matters in which I have testified is attached as Appendix B.

16. Keystone Strategy has provided research support and assistance in my preparation of this report under my supervision, direction, and instruction. My compensation and the compensation of Keystone do not depend on the opinions or testimony that I may give or on the outcome of this case.

II. Assignment

17. On December 16, 2020, a multistate coalition led by the State of Texas filed a lawsuit against Google LLC (Google) asserting violations by Google of federal and state antitrust laws and violations of other state laws, in connection with Google's conduct in the online display advertising industry and as to digital advertising technologies ("Ad Tech" or "Ad Tech stack"). Currently, 16 States (Texas, Alaska, Arkansas, Florida, Idaho, Indiana, Kentucky, Louisiana, Mississippi, Missouri, Montana, Nevada, North Dakota, South Carolina, South Dakota, and Utah) and the Territory of Puerto Rico are Plaintiffs in the case (Plaintiff States). I was retained in February 2024 to provide expert analysis and opinions on behalf of all of the Plaintiff States.

18. I have been asked to address, describe and, where appropriate, provide opinions concerning the following subjects and issues:

- 1) The business, structure, and operation of the ad tech industry, and particularly: (a) the role of publishers, advertisers, and exchanges and other intermediaries in the programmatic buying and selling of website display advertising space; (b) the purpose and function of the multiple

⁷ Chandler-Pepelnjak, John Winston, "Modeling Conversions in Online Advertising" (2010). Graduate Student Theses, Dissertations, & Professional Papers. 670. <https://scholarworks.umt.edu/etd/670>.

⁸ Steele, B., Chandler, J., and, Reddy, S. "Algorithms for Data Science." *Springer International Publishing Switzerland*. 2016. pgs 1-423.

HIGHLY CONFIDENTIAL

technological tools and platforms utilized in the ad tech industry; and (c) the major players in the various segments of that industry, both presently and over time.

- 2) The types of information and data desired, used, or generated by and provided to publishers, advertisers, intermediaries, and owner-operators of ad servers, exchanges, networks, and ad-buying tools to meet their respective operational and economic objectives in the buying and selling of website display advertising space (i.e., "relevant information").
- 3) How tools and relevant information are used by publishers, advertisers, intermediaries, and owner-operators of ad servers, exchanges, networks, and ad-buying tools in the digital marketing and ad tech industries.
- 4) How the possession of, control over, and ability to exclude others from access to relevant information can affect: (a) the buying and selling of website display advertising space; (b) the operational objectives and economic interests of publishers, advertisers, intermediaries, and owner-operators of ad servers, exchanges, networks, and ad buying tools; and (c) the structure of the industry and composition of its players.
- 5) Whether: (a) Google's activities in the ad tech space, involve or have involved the manipulation of and/or failure to disclose information about its auctions, auction rules, or auction mechanics; (b) Google has conflicts of interest in connection with its position in the ad tech industry and digital advertising ecosystem and, if so, whether it has engaged in any self-preferencing conduct in the face of such conflicts; and (c) any Google auction manipulations, failures to disclose, or self-preferencing conduct in the face of conflicts of interest can detrimentally affect, or have negatively impacted, the transparency and overall fairness² of programmatic website display ad auctions.

19. A list of all documents referred to in this report and relied upon by me in forming my opinions in this case is attached as Appendix C. I have reviewed, signed, and complied with the Confidentiality Order entered in this case. My

² Based on my industry experience, I believe a fair and transparent online auction would be built on the principles of equal access to information and a bidding process free from bias or preferential treatment. Bidders should receive the same details regarding the auction's structure, rules, and criteria. Clear communication about the bidding process and how bids will be evaluated is essential to prevent any one participant from gaining an undue advantage. The bidding process itself should be straightforward and open, with all participants fully informed about how to place bids and the timing of these bids. A fair auction should be free from bias or preferential treatment, ensuring that all bids are judged based on the same criteria. This also means avoiding any mechanisms that allow certain participants, such as those with a "last look" advantage, to see other bids before placing their own.

HIGHLY CONFIDENTIAL

supporting team has also read, signed, and complied with the Confidentiality Order entered in this case. I have also reviewed the Stipulation and Order regarding Expert Discovery in this case.

20. I understand that document productions are ongoing in this case and that additional relevant documents may be produced in this case by Google and third parties right before and after I issue this report. I may, and reserve the right to, review and rely on additional documents in conducting my work and forming my opinions in this case. I also reserve the right to supplement or amend this report if my opinions change or require supplementation as a result of my ongoing review of documents.

III. Summary of Opinions

21. In this report, I address the history of advertising and marketing, characteristics of open web display advertising and other advertising channels , the role of data in advertising measurement, and Google's practices and position in these industries.

22. In this report, my primary methodology is the application of my 25 years working in digital marketing including extensive work in every type of ad tech platform discussed within the report. I am applying these years of direct observation and participation in this industry to encapsulate and present the concepts and structures that the vast majority of industry participants would find accurate. I further use my training and experience in the marketing field to ensure that my observations and analyses are consistent with the relevant peer-reviewed literature.

23. My opinions are set forth in detail in this report below and include the following:

- 1) Publishers monetize and advertisers use display advertising for reasons that are specific and different from other forms of advertising generally and other forms of digital advertising specifically.
- 2) Digital marketing is structured into marketing channels, which are used by advertisers in distinct and differentiated ways. Open web display advertising has unique characteristics and serves different purposes or goals for advertisers than other marketing channels.
- 3) A substantial portion of display advertisers purchase their display advertising space through a programmatic auction process, rather than or in addition to guaranteed direct contracts with publishers. Display advertisers use programmatic buying for a variety of reasons, though the primary reason is efficiency. Via programmatic display, advertisers can: (a) increase the number and variety of sites on which their advertisements appear; (b) have greater flexibility to modify or change the types

HIGHLY CONFIDENTIAL

of ads, publishers, and targeted viewers of their ads; and (c) reduce the costs of media buying.

- 4) A substantial portion of publishers offer some portion of their advertising inventory for sale through programmatic auctions. Publishers use programmatic selling for a variety of reasons, but the primary one is yield maximization. Via programmatic selling publishers can: (a) have access to a much wider pool of advertisers, increasing the demand for their advertising space; (b) reduce or eliminate the need for and cost of a direct display ad sales staff; (c) provide a sales channel for remnant display space inventory not sold directly; and (d) maximize the portion of available inventory that is sold. This inventory is subjected to a bidding process and sold to the highest bidding advertiser or to a third party acting on the advertiser's behalf.
- 5) There are a number of types of display auctions, including those with one or more of the following characteristics or structures: (a) first price versus second price; (b) real-time versus one participant having last-look; (c) header bidding versus Google's Open Bidding; and (d) waterfall versus multi-tier versus single-tier. Each of those characteristics is generally understood in the digital advertising and ad tech industries as having a specific algorithmic structure.
- 6) The programmatic purchase and sale of display advertising space is effectuated through services provided by intermediary "ad tech" platforms and tools, which generally include, but are not always limited to: (a) a publisher inventory management system; (b) a publisher ad server; (c) a publisher selling tool; (d) an advertising exchange; (e) an advertiser ad server; and (f) an advertiser buying tool.
- 7) When faced with competitive threats, Google has strategically acquired competitors to maintain and enhance its market position. This approach has enabled Google to eliminate potential rivals and integrate valuable technologies, reinforcing its dominance in the ad tech ecosystem. Through these acquisitions, Google has built its dominant position in the display advertising market.
- 8) Google provides and has provided platforms and tools in each of the foregoing categories, including its DoubleClick for Publishers (DFP) ad server, its Google Ad Exchange (AdX) exchange, its Google Ad Manager (GAM) ad server, and its DV 360 and Google Ads ad buying tools. Google is recognized in display advertising and the ad tech industry as the predominant player in publisher ad servers, ad exchanges, and ad-buying tools.

HIGHLY CONFIDENTIAL

- 9) There are inherent conflicts of interest when a single company provides both sell-side and buy-side platforms and tools, such as publisher ad servers and ad buying tools, respectively. The interests of publishers and advertisers are not generally aligned in a transaction for display advertising space. Conflicts of interest can harm the transparency and fairness of the auction process. These conflicts of interest can be a problem with participants who do not have good alternatives.
- 10) Additional conflicts of interest arise when this single company with both buy-side and sell-side tools is also in the exchange business. Google is such a company, and the digital advertising and ad tech industries generally recognize the existence of Google's multiple conflicts of interest.
- 11) Advertisers and publishers depend on transparency and fairness when they engage in the programmatic website display auction process. These in turn depend in large part upon the nature and extent of the available information regarding that auction and the degree of the participants' and intermediaries' access to necessary information.
- 12) To maximize the cost-effectiveness of their purchase of programmatic display advertising, and to optimize their auction-related strategies and platform choices, advertisers typically need: (a) data and information regarding the mechanics and rules of the auction process; (b) an understanding of the algorithms employed by the intermediary platforms, tools, and exchanges; (c) data and information about the website visitors who will ultimately receive the display ad; (d) information about the space (i.e., impression) where the advertisement will be displayed; (e) information about the commission, share, take rate, mark-up, or other portion of their payment that is retained by intermediaries; and (f) data related to the performance and effectiveness of their ad purchases.
- 13) To maximize their revenues from the programmatic sale of their display advertising space, and to optimize their auction-related strategies and platform choices, particularly price floors, publishers typically need: (a) data and information regarding the mechanics and rules of the auction process; (b) an understanding of the algorithms employed by the intermediary platforms, tools, and exchanges; (c) data and information about the website visitors who will ultimately receive the display ad; (d) information about the commission, share, take rate, mark-up, or other portion of their payments that is retained by the intermediaries; and (e) data related to the performance and effectiveness of their ad sales.
- 14) Restricting or limiting the availability and flow of necessary and critical information to participants and/or intermediaries in the display ad

HIGHLY CONFIDENTIAL

auction process can affect the transparency and overall fairness of the auction process.

- 15) Similarly, denying symmetrical and fair access to inventory, demand, and functionality to some advertisers, publishers, ad servers, exchanges, or ad buying tools involved in an auction (i.e., unequal distribution of information) can harm the transparency and fairness of the auction process.
- 16) Likewise, changing auction rules without adequate or timely disclosure can harm the transparency and fairness of the auction process.
- 17) Google's Bernanke, Global Bernanke, Bell, Reserve Price Optimization (RPO), Dynamic Revenue Sharing (DRS), Poirot, Elmo, Exchange Bidding, Dynamic Allocation (DA), Enhanced Dynamic Allocation (EDA), tying DFP to AdX, and Privacy Sandbox programs and practices entailed one or more of the following: (a) failures to adequately or timely disclose changes to the auction's mechanics and purposes; (b) unwarranted restrictions on material information needed by auction participants and intermediaries; (c) denials of equal and fair access to inventory, demand, and functionality to advertisers, publishers, ad servers, exchanges, or ad buying tools; and (d) conflicts of interest. Those programs and practices jeopardized, and detrimentally affected, transparency and fairness of the auctions in which they were employed.

IV. Origins and Overview of Open Web Display Advertising¹⁰

A. The Advent of Digital Marketing

24. In the past three decades, marketing has evolved significantly. The most pivotal change in the industry happened in the mid-1990s, when marketers began to regard the Internet as a profitable advertising medium. Prior to the digital age, advertising agencies primarily created what is now called “traditional media”: ads of the kind we see on billboards, on television, or in magazines, for example. Most traditional ads are “non-addressable,” which means that they are not targeted at particular individuals. The advent of digital advertising began the era of mass addressability, with marketers extensively targeting ads to individuals.¹¹ In the digital era, marketers also became capable of monitoring in near-real time how individuals respond to ads and online media. As the Internet increased in popularity, capturing more of the attention of individual users, and as advertisers and publishers developed new tools to enhance digital advertising strategies, a new marketing ecosystem emerged.

¹⁰ Web display advertising includes display shown in “walled gardens.” Display advertising outside walled gardens are often referred to as “open web display advertising.”

¹¹ Before the Internet, the only scalable addressable medium was direct mail, which operated on a smaller scale than digital marketing and which was only addressable to the household level.

HIGHLY CONFIDENTIAL

25. The display advertising marketplace trades in the attention of individuals and the information that their online activities reveal. Publishers, like website owners who sell space for ads, want to amass the largest and most valuable audiences so they can earn the highest rates from advertisers.¹² Advertisers, in turn, are able to draw on immense troves of dynamic information about the habits and traits of the audiences they are buying and advertising to. With this information, and with the increasingly sophisticated technologies that both advertisers and publishers develop to access and analyze user information, marketers fine-tune their strategies in order to increase profits.

26. The contemporary display advertising ecosystem is a dynamic, technologically driven, and highly profitable market. To understand how a company like Google rose to prominence in this space, it is helpful to understand how the current ecosystem of display advertising, and the advertising technology (ad tech) industry that supports it, evolved from the early days of digital advertising.

27. As I will elucidate below, advertising is viewed as comprising three key components that I will refer to as the “who, what, and where” of advertising.¹³ The “who” component is the audience¹⁴ to which the advertisement is being shown. In the context of digital marketing, advertisers typically have an incomplete picture of the person who receives an advertisement. The amount of data and type of data available for an individual is a key, defining characteristic of advertising. The “what” of advertising is the ad content, typically called the “creative.”¹⁵ There are hundreds of creative formats online, but the most common are ads that would be recognizable to any regular user of the web: a piece of video, a social media post, or a display banner. The “where” of advertising refers to the context within which the ad is viewed. This context influences the recipient of the ad, but also can radically change the information that advertisers have. For instance, we unfortunately know a great deal about someone searching for “mesothelioma”¹⁶ or “long-term asbestos exposure.” Similarly, social media marketers may leverage highly predictive information about how someone’s friends responded to content on the site.

¹² There may be others.

¹³ Giombi, K., Viator, C., Hoover, J., Tzeng, J., Sullivan, H., Donogue, A., Southwell, B., and Kahwati, L., “The impact of interactive advertising on consumer engagement, recall, and understanding: A scoping systematic review for informing regulatory science.” *PLOS ONE* vol. 2, no. 17. 2022.

¹⁴ Walmsley, B. “Understanding Audiences: A Critical Review of Audience Research.” *Audience Engagement in the Performing Arts*. 2019. pgs. 25-62; See also, Pascucci, F., Savelli, E., Gistri, G., “How digital technologies reshape marketing: evidence from a qualitative investigation” *Italian Journal of Marketing*. 2023. Pgs. 1-32.

¹⁵ Altamira, M., Putri, K., Samdura, R. “The Role of Creative Content in Digital Marketing Strategies in Educational Institution social media (Case Study: Instagram of Vocational Education Program, Universitas Indonesia).” *Proceedings 2022* vol. 83, no. 1. 2023. pgs. 1-12.

¹⁶ Digital Commerce. “The most expensive keywords in paid search, by cost per click and spend” (August 6, 2015). Accessed on June 3, 2024. <https://www.digitalcommerce360.com/2015/08/06/most-expensive-keywords-paid-search-cpc-and-spend/>.

HIGHLY CONFIDENTIAL

28. During the mid-1990s, when advertisers first began to see the Internet as a channel for serving ads to audiences, the first ads were known as banner ads. Banner ads are the kinds of ads that you see at the top of a website. AT&T placed the first clickable banner ad in 1994, and it was quite successful, with a 44% click-through rate. By the late 1990s, however, as the number of Internet users increased, and as banner, and later pop-up ads became less of a novelty, advertisers saw the need to try different channels for accessing responsive users. In this period, search engines began to proliferate, and advertisers began to see search engines as another online channel for serving ads. From there, the online channels for advertising proliferated.

B. Research Sources

29. In this report, I apply my many years of direct observation and participation in this industry to distill down and present the basic concepts and structures that the vast majority of industry participants would find accurate. As I explain the different facets of digital marketing, I will also include sources from the industry discussion online. Marketers are fundamentally practitioners who discuss, explore, and adapt the concepts to in their own public facing media. Moreover, marketing technology changes quickly, so informational websites are often the most reliable resource of new or evolving ideas. In my discussion, I cite several websites in order to reinforce or re-explain my opinions, and to demonstrate the intersubjective agreement across the industry that share my experience and understanding.

30. The sources I cite come in several forms, including those enumerated below:

- 1) Journalistic resources, like *The New York Times* and *The Wall Street Journal*. These kinds of resources adhere to journalistic standards and practices for reporting. There are many journalistic resources that have been covering ad tech for decades, but which may not be familiar to non-practitioners. For example, ClickZ has been covering digital marketing since 1997. TechCrunch has reported on tech and start-ups since 2005. Digiday has closely followed the media industry since 2008 and provides an invaluable source of collective knowledge, in addition to breaking news. Sources like these are considered highly reputable by the industry.
- 2) Corporate websites, like those for HubSpot, Microsoft, and Google, contain important information about the industry. Since these companies have a particular agenda, these sources are used in places where potential bias is unlikely, such as product launch announcements, or where their information conforms to generally accepted industry practices. These websites help to display the shared learning and thinking of companies working in marketing.

HIGHLY CONFIDENTIAL

- 3) Web commentary by industry leaders, like John Battelle, which helps to illustrate experiences that I share with other practitioners.
- 4) Informational websites, like online references or the site of the Interactive Advertising Bureau, which provide definitional support for many of the industry concepts germane here and help explain their evolution. Statista, for example, aggregates data and conducts its own research, and has been providing reliable data to practitioners since 2007.
- 5) Academic papers. Since marketing moves quickly, academic sources cannot always address current iterations of practice. Nevertheless, where academic resources are available for my research and relevant, I turn to these for support.

C. Key Definitions

31. When marketers strategize to place display ads, they take account of many of the following concepts:

- 1) **Impressions.** An impression is the unit of measurement of digital advertising and represents a single instance of a particular ad being shown.¹⁷ One can think of an impression as the viewing of an advertisement. In broadcast TV, there are as many impressions of a single spot as there are televisions turned to that channel when the spot airs. Out-of-home impressions are determined by the number of people passing by the ad at an angle which allows them to see the content. On the web, impressions are views of ads, though there are areas of dispute, such as whether an ad could have actually been seen.¹⁸

Impressions help marketers measure their advertising. Advertisers also take stock of locations where their advertisements are appearing (on which platforms, sites, geographic locations, etc.).¹⁹

- 2) **Reach.** The number of unique people who come into contact with specific marketing content. Reach helps marketers know how well they are saturating an audience with advertising. When consumers engage with advertising, they provide a new, powerful set of data points. Digital marketing practitioners use the metric reach, but typically are actually

¹⁷ The IAB has a technical definition that takes into account things like bot activity, error codes, and when the impression happens. See, IAB. “Glossary of Terminology” (undated). Accessed on June 5, 2024. <https://www.iab.com/insights/glossary-of-terminology/>.

¹⁸ Google Ad Manager Help. “Overview of Viewability and Active View Next: How Active View metrics are calculated” (undated). Accessed on June 4, 2024. <https://support.google.com/admanager/answer/4524488?hl=en>.

¹⁹ When advertising content is shared through online communication and social media, new impressions are generated. Impressions are the currency of most digital advertising.

HIGHLY CONFIDENTIAL

using “cookie reach”—the number of computer browsers reached by a campaign. I will reserve the term “reach” to indicate reach to people and use “cookie reach” when I am describing estimates of reach to computer browsers (e.g., a particular instance of Firefox on someone’s computer).

- 3) **Frequency.** The number of times a specific piece of marketing content came into contact with a single consumer. Impressions, reach, and frequency are related by a straightforward equation: $\text{Impressions} = \text{Reach} \cdot \text{Frequency}$. Typically, in digital marketing, both the reach and frequency figures refer to cookies (i.e., “cookie reach” and “cookie frequency”) and practitioners often elide the distinction. Since reach measures the number of people in contact with specific content, frequency tells us about repetition with the same content. Commonly, marketers seek a sweet spot for frequency—too little and an ad may not make a cognitive impact on consumers; too much and the advertiser will have wasted advertisements that would have been better spent reaching a new audience.

At its most basic, frequency is reported as the ratio of impressions to reach. More sophisticated marketers ask for full frequency distributions. These tell a marketer the percentage of their audience that is reached one time, two times, etc. Regardless of method of reporting and calculation, frequency is typically monitored to ensure that consumers are receiving enough advertising to break through the clutter, but not being “burned out” on the advertisement.²⁰

- 4) **Targeting.** Targeting is a mechanism by which advertisers can narrow the scope of their advertising, seeking homogeneous groups of users to whom they may tailor messages. There is a constant tension between the size of an audience and similarity of its response to advertising offers. Virtually all marketers do some form of targeting. Retargeting, where advertisers show ads to people who have visited their website, typically has the smallest reach but the best performance. There are many other flavors of targeting as we increase the size of the audience, with the most general typically being targeting based on demographics (e.g., women between 35 and 54 years old) or geography (e.g., everyone in the San Francisco area). Many vendors will sell to “interest segments” such as pet owners, travel enthusiasts, or people in the market for a new car.

The following list describes the major types of advertising targeting:

- Demographic Targeting: Involves segmenting the audience based on demographic factors such as age, gender, income, education, and

²⁰ Schmidt, S., and Eisend, M. “Advertising Repetition: A Meta-Analysis on Effective Frequency in Advertising.” *Journal of Advertising* vol. 44, no. 4. Pg. 415-428.

HIGHLY CONFIDENTIAL

occupation. This type of targeting is used to reach consumers who likely have similar needs, desires, or interests based on their demographic profiles.

- Geographic Targeting: Allows advertisers to deliver ads to users based on their geographic location. This can range from broad targeting like country or state to more precise targeting like city or ZIP code. Geographic targeting is particularly useful for local businesses or events.
 - Behavioral Targeting: Based on the browsing behavior, purchase history, and other actions taken by the user. This form of targeting uses collected data to predict which ads might be the most relevant based on past online activity.
 - Contextual Targeting: Ads are shown based on the content of the webpage that the user is viewing. If the content includes certain keywords, ads related to those keywords can be displayed. This ensures that the ads are relevant to the current interests of the viewer.
 - Psychographic Targeting: Focuses on the lifestyles, attitudes, interests, and personality traits of users. This type of targeting attempts to comprehend consumer motivations and preferences on a deeper level, often using data from social media and other online engagements.
 - Retargeting (or remarketing): Involves showing ads to users who have previously visited a particular website or used a specific app but did not complete a purchase or desired action. This is effective in nudging potential customers closer to a purchase.
 - Lookalike Targeting: Targets new users who resemble existing customers in terms of interests, behaviors, and demographics. This is often used to expand reach while maintaining relevancy by tapping into new audiences who are likely to be interested in the advertiser's offerings.²¹
- 5) **Programmatic Advertising:** Programmatic advertising refers to the automated buying and selling of digital ad inventory using software and algorithms. This method leverages data and technology to streamline the ad placement process, targeting specific audiences in real-time and optimizing ad campaigns for better efficiency and performance.

²¹ Tyler, S., Pandey, S., Gabrilovich, E., Josifovski, V. "Retrieval models for audience selection in display advertising." *ResearchGate*. 2011. pgs. 593-598.

HIGHLY CONFIDENTIAL

- 6) **Direct Buying:** Direct Buying in digital advertising refers to the process where advertisers purchase ad inventory directly from publishers or media owners. This approach often involves negotiations and agreements on pricing, placements, and campaign specifics, leading to guaranteed ad placements on selected websites or platforms.
- 7) **Indirect Buying:** Indirect Buying is the process where advertisers purchase ad inventory through intermediaries such as ad networks, ad exchanges, or demand-side platforms (DSPs). This method allows advertisers to reach a broader audience across multiple publishers, often using automated systems to bid for ad placements in real-time.
- 8) **Walled Gardens:** Walled Gardens in digital advertising refer to closed ecosystems where the platform owner controls the ad inventory and data. Prominent examples include Google, Facebook, and Amazon. These platforms provide advertisers with access to their extensive user data and inventory but restrict data sharing outside their ecosystem, creating a "walled" environment. Display advertising outside walled gardens are often referred to as "open web display advertising."
- 9) **CPM:** CPM, or Cost Per Mille, is a pricing model in digital advertising where advertisers pay a set fee for every thousand impressions (views) of their ad.
- 10) **CTR:** CTR, or Click-Through Rate, is a metric that measures the percentage of users who click on an ad after seeing it. It is calculated by dividing the number of clicks by the number of impressions and is expressed as a percentage. CTR is an indicator of an ad's effectiveness and engagement.
- 11) **Website Action:** In digital advertising, a Website Action typically refers to a specific user interaction or behavior that an advertiser aims to track and measure, such as a click, form submission, purchase, or download. Website Actions are critical for evaluating the success of an ad campaign. When I am comparing Website Actions to Conversions, I will typically refer to them as simple "Actions".
- 12) **Conversion:** A Conversion occurs when a user completes a desired Website Action after viewing interacting with an ad. Conversions are used to measure the effectiveness of advertising campaigns in achieving their objectives.
- 13) **Last-Ad Attribution:** Last-Ad Attribution is a model in digital marketing that assigns 100% of the credit for a conversion to the last ad that a user interacted with before making a purchase or completing a desired action. This model assumes that the final touchpoint is the primary

HIGHLY CONFIDENTIAL

driver of the conversion, ignoring the influence of earlier interactions. Last-Ad Attribution is simple to implement and analyze, but it may not fully capture the contribution of other marketing efforts throughout the customer journey.

- 14) **Return on Ad Spend (ROAS) and Return on Investment (ROI):** ROAS measures the revenue generated for every dollar spent on advertising. It is calculated by dividing the total revenue from ad campaigns by the total ad spend. ROAS helps advertisers evaluate the direct financial impact of their ad investments. ROI is a broader financial metric that measures the overall profitability of an investment, including advertising. It is calculated by dividing the net profit by the total investment cost and is expressed as a percentage. ROI provides a comprehensive view of the financial return from marketing efforts. For the purposes of this report these terms will be used interchangeably since the investment in question is advertising spend.
- 15) **Programmatic Guaranteed:** Programmatic Guaranteed is a form of programmatic advertising where ad inventory is reserved in advance at a fixed price. Unlike real-time bidding, this approach guarantees ad placements on specific sites or platforms, combining the automation and efficiency of programmatic buying with the assurance of direct deals.
- 16) **Open Auction:** An Open Auction is a real-time bidding process where multiple advertisers bid for ad inventory in a public auction. Ad exchanges facilitate these auctions, and the highest bidder wins the ad placement. This method allows for broad competition and leads to higher ad prices.
- 17) **Private Auction:** A Private Auction is similar to an open auction but is invitation-only. Publishers invite selected advertisers to participate in the bidding process, often giving them priority access to premium ad inventory. This approach maintains competition while offering more control to the publisher over who bids on their inventory.
- 18) **Header Bidding:** Header Bidding is a programmatic advertising technique where multiple ad exchanges can simultaneously bid on ad inventory before the publisher's ad server makes a call. This method increases competition, potentially leading to higher revenue for publishers by allowing them to receive bids from various sources at once.
- 19) **Waterfall:** The Waterfall model in digital advertising is a sequential ad serving process where inventory is offered to one buyer at a time. If the inventory is not sold, it is passed down to the next network in line, and so on. This method can be less efficient than header bidding due to its


HIGHLY CONFIDENTIAL

linear nature, potentially missing higher bids from lower-priority buyers.

D. The Marketing Funnel

32. One of the foundational conceptual structures in marketing is called the marketing funnel, which helps explain the stages consumers move through as they learn about products for the first time and later commit to them with purchases or loyalty. Also known as the purchasing funnel, or the advertising funnel, the marketing funnel is a classic framework that has been around for decades, first appearing in 1898. It therefore applies both to traditional marketing and digital marketing.

33. In one of its original forms, the marketing funnel, outlines four stages of marketing: awareness, interest, desire, and action (AIDA). Each of these sections describes a stage in the process of appealing to consumers at various stages of the customer journey that leads from product awareness to a purchase or commitment to a product or brand. The framework of the funnel helps marketers understand the different kinds of ads that are appropriate to consumers at different stages of product engagement. Below is an accurate, sample diagram of this funnel:²²



34. Customers do not necessarily move through all the stages of the advertising funnel in a linear or sequential way, and the funnel has evolved considerably since its original formulation to adapt to different kinds of markets and consumers. In general, the funnel helps convey and distinguish key practices in advertising campaigns and how they look in both traditional and digital media.

²² Wikimedia Commons. “The Purchase Funnel” (undated). Accessed on June 3, 2024. <https://commons.wikimedia.org/w/index.php?curid=53843096>.

HIGHLY CONFIDENTIAL

35. The first stage of the funnel is awareness. This is a stage where marketers work to increase brand recognition or awareness to broaden the base of potential consumers. Hence, the awareness face is the largest part of the funnel. At this stage, marketers adopt strategies that are intended to reach large numbers of consumers, typically at a low cost. In this phase, marketers might deploy classic offline advertising like billboards in high-volume locations, for example, or broadcast media, like television ads. Display advertising is well-suited to the stage of the funnel because marketers can reach many people at low cost via this channel. On the other hand, search advertising, being dependent on the behavior of consumers, is uniquely ill-suited for this stage of the funnel.

36. The second stage of the funnel is interest. At this stage, customers are aware of a brand but are looking to learn more about the brand. Therefore, marketers working in this part of the funnel adopt strategies that help to pique consumer curiosity or answer questions about the brand. At this point, content marketing is important, like the information that customers seek on a website or in testimonials. The landing pages of a company website can be very instrumental in this phase of the funnel. Non-branded search can be useful at this stage, as well as targeted display advertising.

37. The third part of the marketing funnel is desire. At this stage, marketers want to create appealing ads that help to heighten desire for products or brands. It is in this phase of marketing where we often see the catchy ads produced by the creative group of a marketing firm. In this stage of marketing, social media marketing and influencer marketing can help build enticing brands. Video can also be quite efficient at this stage of the funnel—video is effective at turning interest into desire, and it becomes a question of targeting to reach the desired audience.

38. The fourth stage of the funnel is action. At this stage, customers convert from advertising to using or embracing a brand or product. Marketers focus on making the brand accessible to consumers who are looking for it, so search marketing, particularly branded search, where marketers deliver ads that are intended to target consumers who are searching for their product online, often play a central role in this stage. Similarly, display retargeting, aimed at consumers who have visited an advertiser's site, can be highly effective here.

39. After someone purchases a product or service, the action phase of the funnel is also a place where marketers acquire information for future retargeting campaigns. Many funnels continue the diagram to include phases of the customer journey that turn a customer into a repeat customer and then a brand loyalist. Email marketing is one of the most effective tools to maintain a relationship with a consumer after they have become a customer.

HIGHLY CONFIDENTIAL

V. Website Display Advertising and Other Digital Marketing Channels

A. Opinions 1–2

40. It is crucial to underscore the distinctiveness of various digital marketing channels from display advertising. Marketers and academics alike recognize these channels—search, social media, digital video, and in-app advertising—as uniquely tailored tools within the marketing landscape, each requiring specific strategies and considerations. This understanding is widely accepted and not subject to significant debate, reflecting the consensus on the need for differentiated approaches to maximize the effectiveness of each channel. As such, it is my opinion that:

41. **Opinion No. 1:** Publishers monetize and advertisers use display advertising for reasons that are specific and different from other forms of advertising generally and other forms of digital advertising specifically

42. **Opinion No. 2:** Digital marketing is structured into marketing channels, which are used by advertisers in distinct and differentiated ways. Open web display advertising has unique characteristics and serves different purposes or goals for advertisers than other marketing channels.

B. Display Advertising

43. Display advertising is generally banner advertising within websites on the Internet. It is typically purchased on a cost-per-thousand-impression (CPM) basis or cost-per-click (CPC) basis.²³ Display advertising is, typically, the easiest way to generate many impressions and reach many people on sites that the advertiser has control over. In the advertising funnel, display likely has the largest share of spend aimed at increasing brand awareness.

44. In display advertising, ads include those designed to retarget or remarket to consumers who have already seen or interacted with previous ads. Whereas traditional display is typically used for brand awareness, retargeted display ads are typically focused on taking a consumer from the “desire” phase to the “action” phase. Retargeted display ads are why, if you put a pair of pants in your shopping cart on the website of a sophisticated marketer, images of those pants haunt your browsing future. Display retargeting is often effective and increases the likelihood that customers will make additional purchases. The reach of display retargeting is typically relatively small, since it is limited by the number of people who have interacted with a website, and its performance is generally effective.

²³ Other purchasing forms include time-based buys, where an advertiser shows ads on a certain placement for a certain length of time and cost-per-acquisition deals.

HIGHLY CONFIDENTIAL

45. Marketers do not consider display advertising to be the same as other marketing channels.²⁴ Display advertising is characterized by its use of visual and multimedia formats to capture audience attention across various websites. Unlike search advertising, which relies on user-initiated queries, display ads proactively reach users based on demographics, interests, or browsing behavior, offering a broader reach that is not limited by search intent. Compared to social media advertising, which integrates seamlessly into the user's social feed and often leverages the platform's rich user data for targeting, display advertising operates across a diverse array of websites. This feature of display enables brands to build widespread awareness and retarget audiences across the web. This capability to visually engage users on a wide range of online spaces, with creative formats like banners and rich media, makes display advertising a unique channel in a marketer's toolkit.

46. The display channel is notable because it can be used at all stages of the funnel. At the top end, advertisers use low-CPM, broad buys reaching millions of people to create awareness about their firm. In the middle of the funnel, we see varying degrees of display targeting, outlined above. At the very bottom of the funnel, we find branded search and display retargeting. For instance, many advertisers set up display retargeting for consumers who have hit the "shopping cart" page but not yet hit the "thank you for purchasing" page.

47. In the early 2000s, advertisers began to see the need to automate the purchase of ad inventory from publishers, and publishers sought to automate the sale of advertising inventory to advertisers. Both groups also wanted to build systems that would allow them to leverage marketing data. The collective effort to create technologies and systems to facilitate this process of automation fostered a sea-change in the process by which display ads are served to consumers.

48. The following example illustrates how programmatic display emerged. Imagine it is the year 2000, and a publisher of a website with space for five display ads on five different pages wishes to sell that space. If those pages were billboards, then you would have five spaces per ad for a total of twenty-five spaces. But when the spaces you have are on a website, the number of ads that you can deliver is determined by the number of visits to your site. Say, for example, you have one thousand visits to your site in a month. This means that instead of having twenty-five spots to sell to advertisers, you have 25,000 ads to sell. The "spots," the places where ads are shown, are typically called "placements" in digital marketing, though terminology can vary. The ads themselves, the 25,000 ad views that our fictive publisher has, are called impressions.

49. Now imagine you are this publisher who wants to sell advertisers 25,000 impressions. An automated system for sales could be helpful since selling via

²⁴ Tsiotsou, R.H., Hatzithomas, L., and Wetzels, M. "Display advertising: the role of context and advertising appeals from a resistance perspective." *Journal of Research in Interactive Marketing* vol. 18, no. 2. 2024. pgs. 198-219.

HIGHLY CONFIDENTIAL

traditional formats requires infrastructure, staff, or both. As the industry strove for efficiency, it arrived at an automated process that allowed vastly more inventory to be sold: an automatic auction. In auctions for ad space (impressions), advertisers can bid on advertising space and if they win the auction, publishers place ads with the help of what is called an advertising server.

50. For many practitioners, the term “programmatic” is synonymous with using auctions to allocate inventory. [REDACTED]

[REDACTED]

[REDACTED]

51. I explain the auction process in greater detail below, but, for now, we can see that programmatic display advertising involves using automated systems that help buyers and sellers bid, buy, sell, and serve (or place) advertisements online.

52. Audience data is also a key part of the programmatic digital advertising space. Consider the example above. When advertisers look for space to serve thousands of ads, they want to find good audiences. To sell ads on a large scale, publishers do not simply sell spots, they sell spots that are seen by individuals and groups of individuals with particular traits and habits. I explore in more detail below the kinds of data that publishers pay attention to negotiate advertising deals. For now, though, we can note that in the programmatic digital space, because of the scale, publishers also automate the process of delivering ads to audiences that advertisers appeal to. The more detail and information advertisers receive about the audiences they are purchasing, the more capable they are of targeting ads.

C. Other Digital Marketing Channels

1. Search Advertising

53. Search advertising was one of the first channels to emerge after the initial era of banner ads, when publishers began to see search engines as a vehicle for serving ads. Search advertising involves placing small text ads on the search results page of search engines such as Google, Bing, or Safari. These ads typically surround what are called the natural search results, which are not paid for.

54. Search ads are different from display ads across all three who-what-where dimensions. In contrast to display ads, search ads are directed at people

[REDACTED]

HIGHLY CONFIDENTIAL

actively seeking something, whereas display ads are intended to captivate and draw consumers toward a product that they are not necessarily seeking. The ads themselves are quite different since search ads are typically text ads while display ads are images. Finally, search ads are delivered in the context of a search results page and so are highly aligned with the context of the search, whereas display ads may or may not be related to the content of the site on which they are displayed.

55. Search advertising distinguishes between non-branded search and branded search. Non-branded search ads are displayed when people search for terms that do not include corporate brands, such as “painkiller.” Non-branded search has been called “the database of intentions” and these ads typically perform extremely well. Reach is limited by the popularity of the search terms the advertiser bids on. Branded search ads are those displayed when someone searches for a branded term, such as “Advil.”

56. The era of search advertising began in the early 2000s and was driven in large part by innovations that Google made to monetize and make profitable search advertising. In the fall of 2000, Google launched AdWords, which made it possible for advertisers to place text ads alongside search results. Through AdWords, advertisers could set how much they were willing to pay per thousand impressions with the higher bidder appearing higher up on the page.

57. By 2002, Google refined AdWords to include a quality score that would determine positions based on a combination of relevance and higher bid. At the time, because Google was the only search engine doing this kind of ranking, it became more popular with advertisers than its rival search engines. Google further developed its search advertising services with AdSense, which launched in 2003. Building on the success of AdWords, Google introduced AdSense to monetize the web more broadly. AdSense allowed website owners to earn revenue by placing Google's ads on their sites. These ads were contextually relevant to the content of the website, meaning that Google's algorithms would analyze the content of a page and serve ads that matched the topics of that content. This relevance increased the likelihood that viewers would click on the ads, benefiting advertisers with higher engagement and publishers with more ad revenue. In short, AdSense allowed publishers to serve content targeted text-based ads from Google AdWords.²⁶

2. Social Media

58. Social media advertising is defined as advertising that takes place within the context of a social media site or app. The actual advertisements on social media often appear like advertisements from other channels such as display, video,

²⁶ Martech. “Google Celebrates 10 Years Of AdSense: Says Over 2 Million Publishers Earned More Than \$7 Billion Last Year” (June 18, 2013). Accessed on June 4, 2024. <https://martech.org/google-celebrates-10-years-of-adsense/#:~:text=Google%20piloted%20its%20content%2Dtargeted,Ad-Sense%20on%20June%202018%2C%20202003>.

HIGHLY CONFIDENTIAL

or content marketing. The difference between social media advertising and these other channels is the use of the data that is unique to social media: the network connections between individuals. Thus, a banner advertisement shown on Facebook would be considered a social media advertisement and not a display advertisement.

59. The marketing under the banner “social media advertising” can take many forms, but I will focus briefly on the kinds of ads that marketers purchase to place on social media platforms (as opposed to other forms of social marketing, which include influencer marketing, content marketing, and earned-media word-of-mouth marketing). Social media ads take the form of ads on sites like Facebook or Instagram or text-like ads on a site like X (formerly Twitter).²⁷

60. Social media marketing is effective, as social media companies, together with the search engines, are the best in the world at predicting what people will be interested in. While search is the “database of intentions”, social media companies make use of the affinity between consumers’ interests and that of their social media connections. Social media companies target different demographics more effectively than channels such as display or search.

61. Several of the most powerful social media platforms launched in the early 2000s. MySpace launched in 2003 and was the first social media platform to reach a global audience. Facebook launched in 2004, and almost immediately began covering its costs with advertising on its platform. By 2006, Facebook was tailoring ads to its users’ demographic traits and interests.

62. When the iPhone launched in 2007, social media channels became an even more powerful platform for serving ads. The same year, Facebook introduced Facebook Ads, which allowed advertisers to connect with users in Facebook’s social network and to target ads to those consumers. When Facebook Ads launched, CEO Mark Zuckerberg said “Facebook Ads represent a completely new way of advertising online. For the last hundred years, media has been pushed out to people, but now marketers are going to be a part of the conversation.”²⁸ Facebook also encouraged businesses to create their own pages which could be a part of user feeds just like individual profiles.

63. After the original launch of Facebook Ads, the platform continued to roll out new ad platforms at regular intervals. In 2008, Facebook created Engagement Ads, which allowed users to engage with ads by commenting, sharing, and becoming fans of the products and services being advertised. In 2011, Facebook introduced


²⁷ It is worthwhile to point out that video sites are sometimes considered social media and sometimes considered video, depending on the context and the features of the site that are being used for advertising.

²⁸ Hub Spot. “How Facebook Ads Have Evolved [+What This Means for Marketers]” (August 24, 2020). Accessed on June 3, 2024. <https://blog.hubspot.com/marketing/history-facebook-adtips-slideshare>.

HIGHLY CONFIDENTIAL

Sponsored Ads, which allowed advertisers to pay to serve ads on Facebook users' news feeds and the news feeds of people in their network, when users interact with the brand or product the advertiser is advertising. The following chart illustrates the many kinds of sponsored ads Facebook's program created:²⁹

Summary: 7 Sponsored Stories types

Story type	Image	Story content	Who sees it
Page Like		Someone liked your Page directly from Facebook or from the Like Box on your website at any point in time.	The friends of your fans.
Page Post		You published a post from your Facebook Page to your fans.	Your current fans.
Page Post Like		One of your fans liked one of your Page posts in the last seven days.	The friends of your fans who liked your Page posts.
App Used and Game Played		Someone used your App or played your Game at least twice or for at least 10 minutes in the last month.	The friends of the people who used your App or played your Game.
App Shared		Someone shared a story from your App in the last seven days.	The friends of the people who shared a story from your App.
Check-in		Someone checked in and/or claimed a deal at one of your claimed Places in the last seven days using Facebook Places.	The friends of the people who checked in or claimed a Deal.
Domain		Someone liked a piece of content on your website using the Like button, shared a piece of content from your website using the Share button, or pasted a link to your website in his status update in the	The friends of the people who liked or shared content from your site.

64. As we can see from the chart above, Facebook Sponsored Ads are display ads that rely heavily on both the activities of individual users and on their social networks to serve advertisements.

65. In 2010, Facebook acquired Instagram, further expanding its reach on social media platforms and making it possible to develop a second platform for advertising revenue. In 2014, Facebook acquired WhatsApp, adding to its extensive trove of user data.

²⁹ TechCrunch. "Summary: 7 Sponsored Stories types" (undated). Accessed on June 5, 2024. <https://techcrunch.com/wp-content/uploads/2011/12/screen-shot-2011-12-20-at-11-37-48-am.png>.

HIGHLY CONFIDENTIAL

66. Facebook's methods for social media advertising continued to proliferate and expand. Between 2018 and 2020 it introduced eight new ways to serve ads on its platforms. It continues to expand and adapt its offerings today.³⁰

67. As Facebook's proliferation of social advertising channels indicates, social media advertising is a channel unto itself distinct from the methods of search and display advertising. Because social media advertising relies extensively on social network transmission, and high levels of consumer engagement, it is a form of advertising that I distinguish from display advertising on the open Internet.

68. It is worth noting in this vein that although Google came to dominate the digital marketing ecosystem in the display advertising space, it did not succeed in the social media space. During the same time period that Facebook proliferated its advertising platform, Google attempted to get a toehold into social media advertising with its platform Google Plus, which it launched in 2011 and shut down in 2019. Google Plus was the fourth time Google attempted to enter the space of social media platforms, yet Google has not been able to gain a toehold in this realm of the advertising market.

3. Digital Video

69. In general, video advertising is an advertising channel that is distinct from social media, search, and display advertising.³¹ There are many dimensions of video advertising that distinguish it from display. Beyond the format differences (images versus videos), video allows marketers to tell a richer story than a banner ad can tell. The opportunity for emotional connection makes video an excellent mid-funnel channel, migrating the general interest in a product to desire. And, to paraphrase Epictetus, desire demands its own attainment. Video is more expensive than display, so it does not make sense as an upper-funnel channel. Creating good video content, while vastly cheaper than a decade ago, is relatively expensive, making it less appropriate for smaller advertisers. Marketers use video ads for campaigns that are capable of paying premium prices for targeted audiences.

70. Video marketing describes the kinds of ads that run before, during, and after videos on sites like YouTube or Twitch. Most advertising on these sites would be considered video advertising and not social media advertising, since most of these campaigns do not make use of network relationships between ad consumers.

71. This channel also includes ads that stand-alone on video platforms—modern infomercials, if you will. Digital video is often thought of as combining the

³⁰ Hub Spot. "How Facebook Ads Have Evolved [+What This Means for Marketers]" (August 24, 2020). Accessed on June 3, 2024. <https://blog.hubspot.com/marketing/history-facebook-adtips-slideshare>.

³¹ Hub Spot. "The Ultimate Guide to Video Marketing" (May 30, 2024). Accessed on June 3, 2024. <https://blog.hubspot.com/marketing/video-marketing>.

HIGHLY CONFIDENTIAL

best of several media. It has the addressability of digital display with the richer storytelling format of TV. This category has exploded in volume of ads and reach as online video has become more popular.

72. While Google was not able to succeed in creating its own social media platform, it acquired one of the largest platforms for ad-supported video—YouTube. YouTube was founded in 2005 and Google bought it in 2006 for \$1.65 billion. In January 2006, YouTube was streaming 25 million video views per day.³² In December 2023, YouTube recorded approximately 98 billion visits on mobile and over eight billion visits from users on desktop devices.³³

73. Another video sharing platform that serves video ads is Twitch. Twitch focuses primarily on live-stream video game and sports content. The scale of Twitch, a subsidiary of Amazon, can be difficult to appreciate for older people. As Business Insider reported in 2018, “962,000 people was the average viewership on Twitch. ... That puts Twitch viewership on par with the likes of MSNBC, CNN, Fox News, and ESPN.”³⁴

4. In-App Advertising

74. In-app advertising is a critical channel for digital marketers seeking to use the intersection of high user engagement and sophisticated targeting capabilities. Like social media advertising, in-app advertising relies on the digital habits of users but does so within the confines of individual applications beyond social media apps, creating a distinct and potent platform for reaching consumers. Context is king for in-app: the defining feature of in-app advertising is that the advertisement is delivered within an application, typically on a smartphone or tablet.

75. In-app advertising represents an evolution in the digital marketing landscape, harnessing the pervasive use of mobile applications to deliver tailored advertisements to users in a manner similar to display. As smartphone adoption skyrocketed in the late 2000s, marketers quickly recognized the untapped potential of mobile apps as a powerful medium for direct consumer engagement. In-app ads are typically integrated within the functionality and design of mobile applications, ranging from subtle banner ads at the bottom of the screen to full-screen video ads that play between levels in a gaming app. The integration allows for a seamless user experience that can be less intrusive than traditional digital advertising methods.

³² Britannica. “YouTube” (undated). Accessed on June 3, 2024. <https://www.britannica.com/topic/YouTube>.

³³ Semrush. "Worldwide visits to YouTube.com from July to December 2023, by device (in billions)." Chart. February 7, 2024. Statista. Accessed June 06, 2024. <https://www.statista.com/statistics/1256720/youtubecom-monthly-visits-by-device/>.

³⁴ Business Insider. “Amazon's streaming service Twitch is pulling in as many viewers as CNN and MSNBC” (February 13, 2018). Accessed on June 3, 2024. <https://www.businessinsider.com/twitch-is-bigger-than-cnn-msnbc-2018-2>.

HIGHLY CONFIDENTIAL

76. In-app marketing shares a common trait with social media advertising—the channel is defined more by where the ads are delivered than the format of the ad. With social media advertising, marketers take advantage of the rich network relationships that exist between consumers. The primary power of in-app advertising lies in its ability to leverage real-time data to optimize ad targeting and that the advertisement is can [REDACTED]

[REDACTED] By accessing mobile device data, advertisers can deliver personalized ads based on user behavior, demographics, and even geographical location. This level of precision ensures that ads are not only seen but are also relevant to the app users, increasing the likelihood of engagement.

77. Moreover, in-app advertising typically offers advertisers robust analytics tools that provide insights into the effectiveness of their ads. Metrics such as click-through rates, engagement times, and conversion rates are available, allowing marketers to optimize their campaigns.

78. The effectiveness of in-app advertising is evident in its ability to maintain high engagement rates. Unlike web browsers, where users can install ad blockers, mobile apps provide a more controlled environment, making it easier for advertisers to capture and retain user attention.

5. Email

79. At its most irritating, email marketing is the spam that clogs one's inbox. The incremental cost of sending an email is vanishingly small, leading to the proliferation of unsolicited emails. Email marketing is an effective tactic and is carried out by most marketers. Its reach is limited by the size of the email lists that the advertiser has compiled, purchased, or both. As I mentioned above, email marketing is often used at the stage of retargeting for consumers who have already made a purchase or for consumers who have interacted already with prior advertising.

80. Marketers consider email a digital marketing channel that is distinct from display advertising. It is not considered a substitute for the channels listed above.

D. Differentiating Website Display from Other Channels Based on Audience, Activity, and Medium

81. As we explore the landscape of digital marketing, it is essential to understand how various channels operate and differ from one another. While display advertising is a fundamental component of the digital marketing ecosystem, each advertising channel has characteristics that set it apart, as I have mentioned above. Display advertising typically involves visual and multimedia ads placed on websites,

HIGHLY CONFIDENTIAL

aiming to capture the audience's attention, build brand awareness, and drive sales. Other channels such as search advertising, social media, digital video, in-app advertising, and email marketing offer distinct approaches and advantages to reaching different types of audiences.

82. In this section, I elucidate the distinctions between these non-display channels with display advertising by examining the three key aspects: the audience (who), the creative (what), and the context (where). By breaking down each channel's attributes and how they differ from display advertising, we can gain a comprehensive understanding of their roles within the broader digital marketing strategy. The following table provides a detailed comparison to highlight these differences and similarities, offering insights into the diverse tactics and methodologies employed across digital marketing channels.

<i>Channel</i>	<i>Who (Audience)</i>	<i>What (Creative)</i>	<i>Where (Context)</i>
<i>Search Advertising</i>	Typically highly intent-driven users searching for specific information	Text-based ads, product listings, tied to keyword and keyword groups	Search engine results pages, influenced by user queries
<i>Social Media</i>	Users with rich demographic and interest data, network connections	Posts for the platform, video ads, sponsored content, images	Social media platforms and feeds, influenced by user engagement
<i>Digital Video</i>	Video platform users, often segmented by viewing habits	Pre-roll, mid-roll, post-roll video ads, stand-alone video ads	Video streaming sites like YouTube and Twitch, surrounding video content
<i>In-App Advertising</i>	Mobile app users, often with behavioral and location data	Banner ads, interstitial ads, rewarded video ads	Within mobile apps, influenced by app usage patterns and real-time data
<i>Email</i>	Subscribers, often segmented by previous interactions and purchases	Text-based emails, HTML emails with multimedia	User email inboxes, personalized based on previous interactions


83. The table clearly illustrates that each digital marketing channel is distinct from display advertising, with unique characteristics in terms of audience,

HIGHLY CONFIDENTIAL

creative formats, and contextual placement. Marketers approach these channels with different strategies and objectives, recognizing that they are not interchangeable with display advertising. For example, search advertising targets intent-driven users with text-based ads, while social media advertising leverages rich demographic data for engaging content within social feeds.

84. Academics also treat these channels separately, as reflected in the scholarly research. Academic studies consistently categorize and analyze search, social media, and video advertising independently of display advertising, highlighting their specific dynamics and effectiveness. There is a consensus among both practitioners and scholars that these distinctions are well-established and not a topic of significant debate. This clear demarcation underscores the importance of tailored strategies, tactics, and analyses for each channel to maximize marketing outcomes.

85. Industry participants similarly treat these channels separately. For instance,



³⁸

E. Differentiating Website Display from Other Channels Based on Price

86. It is exceedingly clear that these channels are thought of as separate and distinct by marketers when it comes to planning and executing their campaigns. It is worth noting that differences in these channels persist through to the pricing of the channels. Marketers are willing to pay different amounts to modify the who-what-where mix.

87. A 2023 study conducted by the media measurement company SEM Rush discusses the CPM in digital advertising across various industries and ad types.³⁹ There are several noteworthy points.

³⁶ Deposition of [REDACTED] 130:6–
130:17. [REDACTED]

³⁷ Deposition of [REDACTED] 83:4–83:9. [REDACTED]

³⁸ Deposition of [REDACTED]
147:20–147:24. [REDACTED]

³⁹ SemRush. “Advertising Trends: CPM Benchmarks by Industry [Study]” (October 23, 2023). Accessed on June 5, 2024. <https://www.semrush.com/blog/advertising-cpm-benchmarks-study/#digital-ad-cpm-by-ad-type>.

HIGHLY CONFIDENTIAL

88. First, regarding video advertising costs: video ads are notably the most expensive format. Mobile video ads specifically have a high CPM, reaching up to \$11.10. In-app and social video ads are also costly, with prices starting at \$9 per thousand impressions.⁴⁰ This indicates that video content, particularly on mobile and social platforms, demands a premium due to potentially higher engagement and effectiveness.

89. Second, regarding display advertising costs, display ads on desktop platforms are the most affordable, with an average CPM of only \$2.50.⁴¹ This lower cost suggests that while desktop display ads may reach a wide audience, they could be less effective or engaging compared to more dynamic video ads on mobile or social channels.


90. Third, regarding industry-specific variations: the rate that advertisers pay for display ads varies significantly across industries. For example, the food delivery industry sees the highest CPM at \$7.63,⁴² indicating a highly competitive market where reaching consumers might cost more. Travel and finance industries also have high CPMs over \$7 and \$6.52 respectively. Conversely, industries like dating and media experience the lowest CPMs, at \$4.44 and \$4.27 respectively, suggesting lower costs for reaching their audiences,⁴³ possibly due to different advertising dynamics or audience engagement levels. These insights reveal the diverse cost structures across different types of digital ads.

⁴⁰ *Id.*

⁴¹ *Id.*

⁴² *Id.*

⁴³ *Id.*

HIGHLY CONFIDENTIAL

91. Advertisers make tradeoffs among these aspects to find pieces of marketing that provide a positive return on investment (ROI) or return on ad spend (ROAS). But these channels, and the types of creative shown on them, are not viewed as equivalent by advertisers, as the discrepancy in CPMs illustrates.

92. Third party testimony further supports differentiation between the various marketing channels. Advertisers seeking to maximize ROI or ROAS will allocate spend across marketing channels, indicating that marketing channels are not [REDACTED]

VI. The Rise of Programmatic Ad Auctions and Placements

A. Opinions 3–5

93. This section covers the rise of programmatic and the basics of auctions for advertising inventory. From the information I detail below, I form the following opinions:

⁴⁴ Deposition of [REDACTED]

98:14–99:6.

[REDACTED]
130:6–130:17.

HIGHLY CONFIDENTIAL

94. **Opinion No. 3:** A substantial portion of display advertisers purchase their display advertising space through a programmatic auction process, rather than or in addition to guaranteed direct contracts with publishers. Display advertisers use programmatic buying for a variety of reasons, though the primary reason is efficiency. Via programmatic display, advertisers can: (a) increase the number and variety of sites on which their advertisements appear; (b) have greater flexibility to modify or change the types of ads, publishers, and targeted viewers of their ads; and (c) reduce the costs of media buying.

95. **Opinion No. 4:** A substantial portion of publishers offer some portion of their advertising inventory for sale through programmatic auctions. Publishers use programmatic selling for a variety of reasons, but the primary one is yield maximization. Via programmatic selling publishers can: (a) have access to a much wider pool of advertisers, increasing the demand for their advertising space; (b) reduce or eliminate the need for and cost of a direct display ad sales staff; (c) provide a sales channel for remnant display space inventory not sold directly; and (d) maximize the portion of available inventory that is sold. This inventory is subjected to a bidding process and sold to the highest bidding advertiser or to a third party acting on the advertiser's behalf.

96. **Opinion No. 5:** There are a number of types of display auctions, including those with one or more of the following characteristics or structures: (a) first price versus second price; (b) real-time versus one participant having last-look; (c) header bidding versus Google's Open Bidding; and (d) waterfall versus multi-tier versus single-tier. Each of those characteristics is generally understood in the digital advertising and ad tech industries as having a specific algorithmic structure.

B. Introduction

97. In 2000, the vast majority of display ad inventory was purchased directly from publishers through human-negotiated deals. Programmatic ad buying, which leverages software and algorithms to automate the process, existed but was still nascent, with only 1% of online advertising spend.⁴⁵

98. The mid-to-late 2000s saw the emergence of ad networks and exchanges that opened up programmatic channels. DSPs gave advertisers the ability to buy across multiple sources in an automated fashion. The advent of real-time bidding (RTB) ad exchanges allowed advertisers to buy impressions through instantaneous auctions. These technological developments set the stage for programmatic's growth.

99. Several factors then coalesced to drive rapid adoption of programmatic display buying over the next decade. Publishers made more of their display inventory

⁴⁵ Choi, H., Melay, C., Balseiro, S. R., & Leary, A. "Online Display Advertising Markets: A Literature Review and Future Directions." *Information Systems Research* vol.31, no. 2. 2020. pgs. iii-vii, 297-652, C2.

HIGHLY CONFIDENTIAL

available through programmatic pipes, providing greater supply. Advertisers were attracted by programmatic's efficiency in reaching targeted audiences across the web at scale. The expansion into video, mobile, and over-the-top (OTT) television inventory transacted programmatically opened new frontiers. Programmatic display ad spending in the U.S. grew from 10% of display in 2010 to 72% of display by 2021⁴⁶, representing a large majority of total display ad spending.

100. According to Statista, programmatic display ad spending in the U.S. grew from \$50 billion in 2018 to a projected \$168 billion by 2024.⁴⁷

101. The forces propelling the rise of programmatic include its ability to combine enhanced audience targeting through data, the efficiency and cost benefits of automation, greater inventory access, and precise attribution measurement. As brands allocated more digital spend programmatically, the channel hit a tipping point and became the new standard for display advertising.

102. The evolution of auctions in digital advertising is a tale of increasing sophistication and technological advancement, reflecting the industry's response to the need for more efficient and effective ad buying processes. Initially, digital ad buying relied heavily on the "waterfall" or "daisy chain" auction model, where ad inventory was offered to buyers sequentially. In this model, if the first buyer in the sequence did not purchase the inventory, it would cascade down to the next buyer, and so forth. This method often resulted in suboptimal fill rates and revenue, as premium inventory might not be sold at its true market value.

103. As technology advanced, the industry moved towards more sophisticated multi-tier auction environments. These setups allowed multiple buyers to bid on the same inventory simultaneously, but in different tiers, based on their perceived priority. This method aimed to increase the chances of inventory being sold at a higher price but could still leave room for inefficiencies, as not all buyers were given equal opportunity at the outset.

104. The shift to single-tier, real-time bidding (RTB) auctions marked a significant evolution. In RTB, all potential buyers bid on inventory in a single auction. This shift, which began in the late 2000s, was driven by advancements in software and technology that enabled the handling of vast amounts of bid data at incredible speeds, ensuring that inventory could be sold at its maximum potential value in a fair and efficient manner.

105. This evolution from waterfall to RTB auctions can be attributed to the increasing complexity of software solutions in the programmatic space, which have

⁴⁶ *Id.*

⁴⁷ Insider Intelligence. "Programmatic digital display advertising spending in the United States from 2018 to 2024 (in billion U.S. dollars)." Chart. February 28, 2023. Statista. Accessed June 06, 2024. <https://www.statista.com/statistics/278727/programmatic-display-ad-spend-in-the-us/>.

HIGHLY CONFIDENTIAL

allowed for more granular data analysis and real-time decision-making. These technological advancements have fundamentally transformed how ad inventory is bought and sold, optimizing revenue for publishers and efficiency for advertisers.

C. Differentiating Programmatic Auctions from Other Forms of Ad Sales

106. In digital advertising, understanding the distinctions between different methods of ad sales is crucial. Programmatic advertising, which includes real-time bidding (RTB) and programmatic direct deals, leverages automated systems and algorithms to buy and sell ad inventory. This method contrasts sharply with traditional, non-programmatic sales approaches that often involve manual negotiations and direct relationships between advertisers and publishers.

107. To understand the different ways that ad sales take place, it is useful to break down ad sales into a 2x2 grid, categorizing them into direct vs. indirect sales and programmatic vs. non-programmatic methods. The types of advertising that fall into each category is summarized in the following table:

	Direct Sales	Indirect Sales
Programmatic	Automated transactions for guaranteed ad placements. Typically, large advertisers and premium publishers.	Real-time bidding (RTB) for ad impressions through ad exchanges. DSPs and SSPs facilitating real-time auctions.
Non-Programmatic	Manual transactions involving direct negotiations and contracts. High-touch sales teams from both advertisers (e.g., luxury brands) and publishers (e.g., magazine websites) engaging in bespoke agreements.	Aggregation and selling of inventory via ad networks. Ad networks aggregating inventory from small publishers and selling to advertisers seeking broader reach.

108. Direct sales typically involve guaranteed ad placements through direct agreements between advertisers and publishers, while indirect sales use intermediaries like DSPs, SSPs, ad networks, and exchanges. Within these categories, programmatic sales introduce automation and real-time processing, enhancing efficiency and targeting precision, whereas non-programmatic sales rely on manual processes and predefined agreements. This section will explore these differences, highlighting how each method operates and the implications for marketers.

HIGHLY CONFIDENTIAL

109. Ad sales can be broadly categorized into direct and indirect sales, each with programmatic and non-programmatic methods. The difference lies in the process used to execute the transaction for the inventory.

110. Direct sales involve direct negotiations between advertisers and publishers, ensuring premium placements and guaranteed impressions. Traditional direct sales rely on manual processes, contracts, and personalized agreements, emphasizing relationship-based transactions. Conversely, programmatic direct, or automated guaranteed, uses automated systems to facilitate these direct deals, maintaining the pre-negotiated terms and guaranteed impressions but adding the efficiency of automation.

111. Indirect sales focus on selling ad inventory through intermediaries, such as ad networks or exchanges, typically for remnant inventory. Remnant inventory denotes ad impressions that have not been sold via guaranteed deals. Traditional indirect sales involve ad networks that aggregate inventory from multiple publishers, using less automation and often requiring human negotiation. Programmatic indirect sales, however, leverage real-time bidding (RTB) through ad exchanges, employing algorithms and automated systems to buy and sell ad inventory in real-time based on instantaneous auctions. This method significantly enhances efficiency and targeting precision.

112. If we turn our attention to programmatic versus non-programmatic advertising, we see a split primarily along the lines of efficiency.

113. Programmatic advertising revolutionizes ad sales by using automated systems and algorithms to facilitate buying and selling. Real-time bidding (RTB) is a key component, allowing advertisers to bid on individual ad impressions in real-time auctions, enhancing flexibility and efficiency. Programmatic direct combines automation with pre-negotiated deals, ensuring guaranteed impressions with the ease of automated processing.

114. Non-programmatic advertising, on the other hand, relies on manual transactions and traditional ad networks. Manual transactions involve human negotiation, insertion orders, and personalized agreements, making the process slower and more labor-intensive. Ad networks aggregate inventory from various publishers, but the process lacks the real-time capabilities of programmatic methods. This approach often results in less precise targeting and reduced efficiency compared to programmatic advertising.

D. The Advantages and Basic Rules of Programmatic Auctions

115. While ad exchanges are the marketplace where advertisers and publishers come together, and with the help of their ad buying tools and supply side platforms, buy and sell ads, the auction is the central activity through which ads are bought and sold on exchanges.

HIGHLY CONFIDENTIAL

116. The structure and rules of auctions are important to understand to grasp industry norms in the realm of exchanges and in the ad tech industry in general. The following explains the elements and basic workings of an ad auction.

117. Ad exchanges are where advertisers come together with publishers, with advertisers buying and publishers selling ad impressions through real-time auctions. Publishers use publisher ad servers to facilitate the selling of their inventory. Section VII explains the structure of the ad tech ecosystem in detail.

118. For publishers, being able to make available general and remnant inventory to multiple bidders to hundreds of potential buyers present significant advantages. The auction environment can also allow publishers to earn higher bids for inventory that might otherwise be sold to a narrower group of advertisers.

119. For advertisers, the auction environment makes available advertising space across hundreds of platforms that can be bought and served almost instantaneously, so it makes the greatest possible efficiency possible. Because of the broad range of inventory available, the auction environment can also give advertisers access to inventory that meets their optimization criteria at the lowest cost. [REDACTED]
[REDACTED] that the rise of exchanges accompanied a performance improvement:



48

120. Given the benefits to buyers and sellers, it is not surprising that since the first ad auctions in 2007 and 2008, the inventory sold on ad exchanges has increased dramatically. So too has the number of entities buying and selling ads. At present, billions of transactions in advertising sales occur each day.⁴⁹

121. One of the ways that publishers are able to provide more and more of their inventory to advertisers is through header bidding. This is when publishers offer their inventory to many different advertising exchanges before making calls to their ad servers. Header bidding is described lucidly by the site clearcode.cc, who offer an explanation that matches my own understanding, as a kind of advanced bidding “that enables publishers to simultaneously collect multiple bids from a number of demand sources (not only from their ad server) on all of their ad inventory prior to a

⁴⁸ Deposition of [REDACTED] 52:10-55:13, [REDACTED]

⁴⁹ [REDACTED] 107:4-16, 20-24. Usage of Google’s exchange, AdX, has grown as well. See GOOG-AT-MDL-000016711 at -735.

HIGHLY CONFIDENTIAL

sale.”⁵⁰ Clearcode.cc also explains how header bidding gives publishers more information about bids:

[Header bidding] allows publishers to “see” which demand sources (e.g., demand-side platforms) and ad networks connected to ad exchanges and supply-side platforms [SSPs]) are placing the bids and the monetary value of their bids (e.g., \$1.05), allowing them to get the highest cost per mille (CPM) possible. As bids are received before a publisher’s ad server is called, they are able to compete with a publisher’s premium deals, meaning advertisers have a better chance of winning the impression and publishers earn more money for their inventory.⁵¹

^{122.} As a result of header bidding, many publishers have increased their profits by forty percent.⁵²

123. In response to the rise of header bidding and its impact on Google's ad tech dominance, Google launched Open Bidding (originally known as Exchange Bidding). Open Bidding was created as an alternative to header bidding that operated server-side rather than in the user's browser. A server-side auction means that the auction takes place on Google's servers, where multiple ad exchanges and demand-side platforms (DSPs) would bid on the publisher's inventory in real-time. Open Bidding aimed to offer similar auction-dynamic features as header bidding, such as increased competition and higher revenue for publishers. Google launched the beta for Open Bidding on June 8, 2017.⁵³

124. Critics and competitors worried that Open Bidding was designed to undermine the advantages of header bidding and return control over the ad tech ecosystem to Google. For instance, [REDACTED]

said that, at the time:

55

⁵⁰ Clear Code. "What is Header Bidding and How Does it Work?" (April 16, 2024). Accessed on June 2024. <https://clearcode.cc/blog/what-is-header-bidding>.

51 *Id.*

⁵² Search Engine Journal. "What Is Header Bidding And How To Implement It" (undated). Accessed on June 4, 2024. <https://www.searchenginejournal.com/header-bidding/389013/>.

⁵³ Google Ad Manager. "Announcing Exchange Bidding open beta" (June 8, 2017). Accessed on June 6, 2024. <https://blog.google/products/admanager/announcing-exchange-bidding-open-beta/>.

54

55

HIGHLY CONFIDENTIAL

125. It is easy to understand those worries. Since Open Bidding is conducted on the auction server-side, Google maintains greater control over the bidding process and can potentially prioritize its own exchange over others. Additionally, Open Bidding requires publishers to use Google's ad server, further entrenching Google's dominance and limiting the opportunities for independent ad tech companies to compete. While Open Bidding initially offered certain technical benefits, such as reduced latency and simplified implementation, it raised and continues to raise concerns about transparency and fair competition in the programmatic advertising market.

126. [REDACTED]. Within Google, header bidding was seen as a threat to Google's exchange hegemony. [REDACTED]

[REDACTED]⁵⁶ wrote November 2, 2016: "There was an understanding of the existential threat posed by header bidding, especially HB wrappers⁵⁷ from [REDACTED], etc.⁵⁸"

127. [REDACTED], the [REDACTED] [REDACTED], crystalizes exactly what is being threatened by header bidding: "Not clear to me how our strategic place in the infra helps us if FB can just do header bidding and get around us."⁵⁹ Google's exchange dominance, coupled with its dominant positions on the buy side and sell side, allowed Google to create a world where no one could "get around them." And a tollbooth that no one can get around can be a lucrative one.

128. At the auction stage, there are two main kinds: first price and second price. In a first price auction, the highest bidder wins the ad inventory up for bid.

129. In a second price auction, the winning bid pays \$0.01 more than the second-highest bid, as in this image:⁶⁰

⁵⁶ The titles in tech can be opaque. This title refers to a senior role responsible for managing and promoting the tools and services that advertisers use to buy ad space through Google's advertising exchange, known as DoubleClick AdX. This role likely involves overseeing the business aspects of Google's ad-buying services, focusing on increasing their usage and revenue.

⁵⁷ HB Wrappers, or Header Bidding Wrappers, are pieces of software used by publishers to manage multiple header bidding partners efficiently. They provide a framework for running header bidding auctions, where multiple demand sources bid on the same ad inventory simultaneously.

⁵⁸ This list refers to competitors who provide tools to publishers. The "etc." probably includes companies such as Rubicon Project, AppNexus, and Criteo. Header Bidding is a threat to Google because it would allow these competitors to bypass Google's dominance as an exchange. These companies make up the short list of companies who have the ads expertise and the software resources to build enterprise-class HB Wrappers. With this sort of tool, publishers could tap into a diverse range of demand sources, increasing the competition for their ad space and potentially driving up their ad revenues. This migration from Google to a competitor is the threat.

⁵⁹ GOOG-TEX-00090151. "Re: Jedi++ Go to Market" (November 2, 2016). [REDACTED] among others.

⁶⁰ GOOG-TEX-00110540. "Re: Facebook / FAN announcement" (March 24, 2017). [REDACTED]

⁶¹ Kevel. "What are Ad Auctions? The Definitive Guide" (May 16, 2024). Accessed on June 4, 2024. <https://www.kevel.com/blog/ad-auctions>.

HIGHLY CONFIDENTIAL



130. Second-price auctions with sealed bids, may seem strange, but they have unique benefits for advertisers and publishers alike.⁶² The key benefit is that, according to auction theory, sealed-bid second-price auctions allow bidders to bid the true value of an impression. This saves advertisers time and effort, since they do not need to try to game the auction by estimating competitive bids, uncover dynamic price floors, engage in “bid shading,” or model auction dynamics.

131. Since advertisers can bid their true estimated worth of the impression, publishers benefit as well. They are likely to enjoy greater revenue since advertisers are freed from practices like bid shading. The ease and simplicity of the method can encourage additional participation by advertisers, ultimately driving up the value of the inventory.

132. There are two particular flavors of auctions that are germane to my report: real-time bidding (RTB) auctions and auctions where someone has a “last look”. As mentioned above, RTB auctions are those were all bids are “opened” at the same time and a winner is determined. An auction with a last-look participant operates differently.

133. "Last look" refers to a scenario where one participant is given the final opportunity to bid on an ad impression after all other bids have been submitted. This privileged position allows the entity with the "last look" to see all the competing bids and then decide whether to outbid the highest bid or let it stand.

134. The "last look" advantage is significant for several reasons. First, it provides a competitive edge, enabling the privileged participant to always submit the highest bid and consistently win the auction for valuable ad impressions. This

⁶² Vickrey, W. (1961). “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *Journal of Finance*, 16, 123-145. Source, 5Vickrey, W. “Counterspeculation, Auctions, And Competitive Sealed Tenders” *The Journal of Finance* vol. 16. no. 1. March 1961. <https://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1961.tb02789.x>.

HIGHLY CONFIDENTIAL

competitive edge leads to a higher win rate compared to other bidders who do not have the opportunity to see competing bids before submitting their own.

135. Second, with visibility of all other bids, the "last look" participant can strategically bid just above the highest competing bid, minimizing their costs while still securing the impression. This optimized bidding allows them to maximize their return on investment by paying only slightly more than the next highest bidder, rather than potentially overbidding in a blind auction.

136. Third, the "last look" privilege provides valuable insights into competitors' bidding strategies and market prices. This data can be used to refine their own bidding algorithms, better understand market dynamics, and improve future bidding decisions.

137. Lastly, entities with the "last look" can exert significant influence over the market by consistently winning key ad impressions. This leads to a consolidation of market power, making it more difficult for smaller or less privileged competitors to compete effectively.

138. "Last look" raises concerns about fairness and competition in the programmatic advertising ecosystem. The ability to consistently outbid competitors creates an uneven playing field, where the privileged participant dominates auctions and limit opportunities for other bidders. This leads to reduced competition, higher costs for advertisers, and potentially lower revenues for publishers who may not receive the full value of their ad inventory due to the strategic underbidding facilitated by the "last look" advantage.

139. Once an auction takes place, almost instantaneously, the ad server will place an advertisement automatically in the location and with the specifications that the advertiser has purchased in buying the ad. It is helpful to remember that once the auction takes place and the ad is being served, not just the publisher and advertiser are involved with the transaction, but now also the consumer.

VII. The Ad Tech Ecosystem

A. Opinion 6

140. The programmatic purchase and sale of display advertising space involves a complex ecosystem of intermediary "ad tech" platforms and tools that streamline and optimize the process. These intermediaries play crucial roles in managing inventory, serving ads, and facilitating transactions between buyers and sellers. This system includes several key components: a publisher inventory management system, a publisher ad server, an advertising exchange, an advertiser ad server, and an advertiser buying tool. Each of these elements works in concert to ensure the efficient and effective delivery of digital ads across various platforms. In light of my discussion of the ad tech ecosystem, I offer the following opinion:

HIGHLY CONFIDENTIAL

141. **Opinion No. 6:** The programmatic purchase and sale of display advertising space is effectuated through services provided by intermediary “ad tech” platforms and tools, which generally include, but are not always limited to: (a) a publisher inventory management system; (b) a publisher ad server; (c) a publisher selling tool; (d) an advertising exchange; (e) an advertiser ad server; and (f) an advertiser buying tool.

B. Introduction

142. As mentioned in the introduction to digital marketing, advertising comprises three components: placements, creatives, and audiences. These components can be thought of in the following way: the creative content of the ad, also known as the ad creative (the “what”), the placement of the ad (the “where”), and the audience or individual to whom the ad is served (the “who”). This section is devoted to ad tech, which is the “how” of digital advertising.

143. In general, one can think of the ad tech ecosystem as the complex system of entities that facilitate digital advertising. At its most basic level, this system can be understood as a system driven by supply and demand. On the supply side, there are publishers, like websites with space where ads can be displayed. On the demand side, there are advertisers who seek to display their ads to a specific audience. In order to scale up and facilitate the transactions between the two sides, a broader set of entities have emerged, which includes supply side platforms (SSPs) which help aggregate advertising inventory, demand side platforms (DSPs) which help aggregate and place advertising inventory, and exchanges which facilitate the exchange between the two sides. Each entity in this ecosystem faces challenges unique to the task that they aim to complete. Following the path of each entity through this system, it is clearer how these challenges become more or less difficult to solve based on the actions of other entities in the system.

144. As discussed below, there are many different entities in ad tech. These entities, however, have not remained fixed over time. One reason for this lack of fixity is product extension. For instance, at one point publisher ad servers, publisher inventory management companies, yield optimization companies, and SSPs were highly distinct market categories. I discuss the ad serving process below. Inventory management systems once were stand-alone systems that helped publishers forecast their upcoming inventory, keep track of what inventory had been sold, and make sure the publisher met its contracted delivery targets. Yield optimization companies provided services to help publishers find the optimal price for their inventory.

145. For instance, Microsoft acquired AdECN⁶³, an ad exchange platform, with the goal of bolstering its yield optimization offerings. Similarly, Google acquired

⁶³ Microsoft. “Microsoft to Acquire AdECN, Inc” (July 26, 2007). Accessed on June 4, 2024.
<https://news.microsoft.com/2007/07/26/microsoft-to-acquire-adecn-inc/>.

HIGHLY CONFIDENTIAL

AdMeld in December of 2011⁶⁴ to improve this facet of the business. The distinction between these types of companies have evaporated in the last 15 or so years, with all of these products and product features now being subsumed under the category of publisher ad servers. In the sections below, I explain the historical origins of the tools and point out the contraction in roles that have occurred over time.

146. It may be helpful to imagine the process of ad-serving as an interaction between three computers. One computer belongs to the user and holds the user's web browser, which will ultimately display the rendered webpage with any ads on that page. The second computer is the publisher's web server. The publisher has a website with space for ads. The user pings the publisher's web server. Then the publisher's web server goes back and forth with the user and a third computer, the advertiser's ad server. The ad server has all of the ad creative files and impressions are counted.⁶⁵ Eventually the publisher and ad servers leave with data, and the user sees an ad on the website. This happens in seconds.

147. Outlined below is a more detailed, step by step account of the process⁶⁶:

- 1) User requests a webpage: Imagine you're on the Internet, and you want to visit a website. You either type the website's address into your browser or click on a link.
- 2) Publisher's server receives the request and prepares to serve the webpage: When you want to see a webpage, your request goes to a big computer called a server that holds that webpage. This server gets ready to send you the webpage you asked for.
- 3) Publisher's server sends an ad request along with user data to the ad server: Along with the webpage you want to see, the server also sends a message to another computer called an ad server. This message includes some information about you, like your location or interests, to help show you ads that you might be interested in.
- 4) Ad server receives the request and processes it, selecting suitable ads: The Ad Server gets the message and decides which ads would be best to show you. It looks at things like what you might like and what ads are available.


⁶⁴ Google Official Blog. "Take a walk on the sell-side" (December 2, 2011). Accessed on June 7, 2024. <https://googleblog.blogspot.com/2011/12/take-walk-on-sell-side.html> See also, Hong, A., Bhattacharyya, D., & Geis, G. (2012). "The Role of M&A in Market Convergence: Amazon, Apple, Google and Microsoft." *Proceedings of 18th International Business Research Conference 2012*. 2012.

⁶⁵ Interactive Audience Measurement. "Interactive Audience Measurement and Advertising Campaign Reporting and Audit Guidelines" (September 2004). Accessed on June 4, 2024.

⁶⁶ ClearCode. "Ad Serving" (undated). Accessed on June 6, 2024. <https://adtechbook.clearcode.cc/ad-serving/>.

HIGHLY CONFIDENTIAL

- 5) Ad selection process: The ad server picks out the ads that it thinks you'll find interesting from all the options it has.
- 6) Ad server requests advertiser's ad tag: Once it knows which ads to show, the server asks the companies that made those ads to send them over.
- 7) Advertiser's server sends the ad tag: The companies that made the ads get the message and send over the details about their ads, like how they should look and where they should go.
- 8) Ad creative request: The ad server then asks for the actual images or videos that make up the ads from the companies.
- 9) Advertiser's server sends the ad creative: The companies send over the pictures or videos, and the ad server gets them ready to show to you.
- 10) Advertiser's server redirects to a 1x1 pixel: After showing you the ad, the companies may send a little invisible pixel to your browser. It's like a code that helps them know if you saw the ad.
- 11) 1x1 Pixel Request: Your browser gets the little pixel, and it tells the companies that you saw the ad.
- 12) Advertiser's server records impression and counts it: The companies know that you saw their ad because of the little pixel, and they keep track of it.
- 13) Advertiser's ad renders: Finally, the ad shows up on the webpage you wanted to see. It might be a picture or a video, but either way, it's there for you to look at. This is one of the main ways that ads end up on the websites you visit.

HIGHLY CONFIDENTIAL

C. Publisher Ad Servers

148. Publishers create and sell space for advertising. To sell ad space, publishers begin by creating an online presence that attracts audiences. They also track how audiences respond to their content so that they can provide information to potential advertisers about the size and characteristics of the consumers they attract.

149. There are several ways that publishers attract audiences, which, in turn, can be attractive to advertisers. In general, publishers want to create large audiences so that they can sell more ads. As I explained above, display ads are typically sold on a CPM (cost-per-thousand-impressions) basis. The larger the audience, the more impressions a publisher can sell. All else being equal, the more impressions a publisher can sell, the more money they can make. Consider the sports news site

HIGHLY CONFIDENTIAL

ESPN, for example, which is a major publisher of ads. Because of its extremely large user base, ESPN is often a highly attractive publisher for potential advertisers.

150. In addition to amassing audience attention, publishers are also interested in acquiring information about audience members so that inventory can be repackaged, typically using data. This repackaging allows publishers to increase the value of their inventory, since advertisers can use this information to target their advertising campaigns at particular kinds of audiences. Again, a platform like ESPN not only commands a large audience, but it also possesses extensive and highly refined data about the media consumption habits of its users. It therefore can offer advertisers not only large audiences, but also information about individual demographics, and audience patterns. In so doing, a publisher like ESPN can help advertisers target their audiences more effectively. For instance, ESPN could sell audience segments such as “people who follow two or more Texas professional sports teams” or “people whose favorite NBA team has just won a game” or “avid soccer fans.”

151. Two additional characteristics of audiences that publishers highlight to attract advertisers are concentration and value. Some advertisers seek audiences that are focused on specific products or services. Consider a website like Investopedia, for example, which attracts consumers interested in learning about investing. This kind of audience is focused on a specific kind of product. So, it could be said that it is a highly concentrated audience. This audience could also be seen as potentially high value, since consumers looking at investment sites often have more expendable income, and so may be more likely to consume products advertised on that site.

152. For publishers to provide ad space to advertisers, they rely on their publisher ad server. At one point these were third-party tools that helped manage the advertising operations, or “ad ops,” of the business. Two key technologies, inventory management systems and ad servers, helped publishers make their inventory available to advertisers, place ads in real time, and participate in auctions. These technologies had significant overlap and it was difficult to understand the differences between them. Now these facets of the business of publishing are unified under publisher ad servers.

153. Additionally, some publishers, especially those working at a large scale, used supply side platforms (SSPs) which allowed publishers to connect their inventory to large multiple exchanges, DSPs, or networks at once. They helped publishers sell inventory that might otherwise have gone unsold. By offering their inventory to the widest possible range of buyers, SSPs helped publishers maximize profits. Again, this distinction between SSPs and publisher tools has largely vanished in the current digital marketing ecosystem. Vendors that currently sell publisher ad server

HIGHLY CONFIDENTIAL

technology include Google, OpenX, PubMatic, Rubicon Project, AppNexus, Right Media and AOL.⁶⁷

154. Publisher inventory is perishable—once someone navigates away from a web page, those impressions are lost. Large publishers typically have a vast amount of this perishable inventory. Further down, I will go into detail about the implication of this perishable-but-plentiful inventory reality on the supply side. For now, we can note that when publishers possess large inventory, what they are offering to advertisers is not only space for ads, but space for ads that will appear in real time before consumers with specific traits.

155. As mentioned above, advertising is about placements, creatives and audiences, the where-what-who of advertising. Publishers provide the placement and the audience, while advertisers tend to provide the creative. Most impressions are provided by publishers who operate at a large scale, with millions or billions of impressions available every month, and so publishers often need scalable systems that can be used to help advertisers find the specific kinds of placements and audiences they want.

156. One of the services that publisher ad servers provide is inventory management. Publishers use inventory management systems to help them ensure that they satisfy the contractual obligations for the inventory they have sold. For example, ESPN may have sold Nike 30 million impressions in a given month. Nike will be dissatisfied if those impressions are clumped into just a couple of days, so ESPN must strive for even allocation of these impressions with about one million per day. For every individual impression, though, there may be dozens of advertisers who could have a claim on the impression based on the various types of targeted and untargeted deals ESPN struck. The inventory management system helps the publisher fulfill these obligations and minimize mistakes in publishing.

157. In addition to managing their inventory, publishers use publisher ad servers to execute the mechanical ad serving aspect of their business. This portion of the tech stack comprises specialized platforms that facilitate the delivery and management of advertisements on digital media properties such as websites and apps. These servers are responsible for determining which specific advertisements to display to users based on a variety of factors, including the advertiser's targeting criteria, the content of the website, and user data such as demographics and browsing history. Additionally, publisher ad servers track ad performance, providing vital metrics like impressions, clicks, and conversion rates to optimize ad placement and effectiveness.

158. One of the most widely used and comprehensive ad servers is Google's DoubleClick for Publishers (DFP). DFP emerged when Google bought the ad server

⁶⁷ Digiday. "WTF is a supply-side platform" (January 22, 2014). Accessed on June 4, 2024. <https://digiday.com/media/wtf-supply-side-platform/>.

HIGHLY CONFIDENTIAL

DoubleClick in 2008. Google has rebranded this entity, naming the server Google Ad Manager (GAM). But most marketers regard DFP as one of the keystone servers in the ad tech industry.

159. It is rare for publishers to build a custom ad server due to the complexity, cost, and maintenance involved. The core of a publisher's business is creating content and building audience, typically on relatively thin margins. Building an ad server is an undertaking that costs scores of millions of dollars, which is beyond the reach of most publishers. Instead, most opt for third-party ad servers, which provide ready-to-use, customizable solutions that can be relatively quickly deployed and can be maintained with fewer resources. Third-party providers like Google Ad Manager (GAM), Kevel, and Amazon Ad Server offer robust features that include targeting, ad delivery, and reporting capabilities, which are essential for effective ad management and optimization. GAM is the largest of these ad servers. Similarly, it is rare for publishers to use multiple ad servers to serve ads in the same marketing channel. One of these reasons is the complexity of these systems. There are no "economies of scale" in setting up multiple ad servers. The more important reason relates to one of the critical pieces of functionality for a publisher ad server: yield optimization. Yield optimization is the name given to the process publishers undergo to find optimal prices for their inventory. The statistical models underpinning the yield optimization process require comprehensive data to perform accurately. It is more common for publishers to use multiple ad servers across channels, since the process of serving, say, video ads is quite different at a technical level from serving display ads. Similarly, the metrics to analyze the performance of these campaigns can be quite different.

160. Publishers are acutely aware that the quality of advertisements displayed on their sites directly impacts user experience and, consequently, their brand reputation. To ensure that the ads align with their audience's preferences and maintain a high standard of quality, publishers use ad monitoring systems. These systems not only track the performance of ads in terms of engagement and revenue but also scrutinize the content of the advertisements to prevent the display of inappropriate or intrusive ads.

161. Moreover, many publishers set strict ad policies that advertisers must comply with, which cover aspects such as ad format, content appropriateness, and the level of animation or sound in an ad. Publishers are vigilant to safeguard the user experience by ensuring that the ads are not only relevant and engaging but also respectful of the user's online environment.

162. For advertiser ad buying tools, as I discuss below, there is a bifurcation in the market, with some tools being used by large advertisers and some being used by small advertisers. This split is not seen in the publisher ad server market. Most publisher tools, such as Google Ad Manager, are designed to be scalable and versatile, serving a broad range of publishers regardless of size. These tools offer features like ad serving, inventory management, and yield optimization, which are essential for

HIGHLY CONFIDENTIAL

both small and large publishers. Google's tool for large advertisers, DV360, has a minimum spend requirement.⁶⁸

163. While there are certain features and services within these tools that might be more beneficial or accessible to larger publishers—such as advanced analytics, custom integrations, and higher levels of support—the core functionalities remain the same across the board. The primary goal for publisher tools is to maximize ad revenue and streamline ad operations, objectives that are shared by publishers of all sizes.

164. This lack of a clear split could be due to the nature of the publisher ecosystem, where even small publishers can benefit from comprehensive ad management solutions, and large publishers require scalable tools that can handle their extensive inventory and complex operations. The versatility and scalability of publisher tools are designed to accommodate this wide range of needs within a single platform.

165. The lack of a clear split in publisher tools between small and large publishers is also influenced by the underlying technology. At its core, implementing ad management and serving solutions involves embedding JavaScript tags into web pages, a process that is relatively straightforward and can be executed by publishers of all sizes. This fundamental simplicity ensures that even small publishers can access and utilize advanced ad management systems without extensive technical expertise or resources.

166. However, while the initial implementation may be simple, the complexities increase with scale and specific requirements. Larger publishers often demand more sophisticated integrations and features to handle their extensive inventory and high traffic volumes. For example, large-scale publishers may require custom implementations to integrate seamlessly with their existing content management systems, data management platforms, and other tools. They also need robust infrastructure to manage high volumes of traffic and ad impressions, ensuring efficient and effective ad serving with minimal latency. Any tool capable of supporting a large advertiser will easily accommodate the needs of a small advertiser as well.

⁶⁸ Programmads, “Why use DV360 instead of Google Ads?” <https://programmads.com/project/why-use-display-video-360-instead-of-google-ads/>. Accessed May 23, 2024; Google does not publicly disclose pricing for DV360, but several online sources list \$50,000 as the minimum monthly spend. Ganz E., ADCORE Blog, “What is DV360 and How to Start Advertising,” (March 11, 2024) <https://www.adcore.com/blog/what-is-dv360-and-how-to-start-advertising>. Accessed May 23, 2024. The Google Marketing Platform page places DV360 under the “enterprise” tab. <https://marketing-platform.google.com/about/display-video-360/> Accessed May 23, 2024; DV360 used to be called “DoubleClick Bid Manager” (DBM) and is sometimes referred to internally at Google as “XBid”. Internal documents often refer to DV360 using these terms. Display & Video 360 Help, “Introducing Google Marketing Platform,”, <https://support.google.com/displayvideo/answer/9015629>. Accessed June 4, 2024. Nitish Korula Deposition, (April 19, 2024) at 118:11-13.

HIGHLY CONFIDENTIAL

D. Advertiser Ad Buying Tools

1. Overview

167. On the demand side, we have advertisers. Advertisers are typically playing a game that is a trade-off between the audience's size and quality. For example, imagine you are the investment bank Morgan Stanley trying to get people to put ten thousand dollars into a mutual fund. At the moment, there may only be a few hundred people in the United States who are prepared and inclined at the moment to invest ten thousand dollars in mutual funds. One option would be to run an advertising campaign that advertised broadly, hoping to find the few people prepared to invest that amount today. The cost of a campaign increases with the number of impressions, but even at high spend levels, a broad campaign may not reach this very specific subset of people. Another option would be to try to create an advertising campaign to target these specific people. That campaign would have a very small, but highly valuable audience. As a marketer, you might therefore be willing to pay a high price to advertise to that audience, with premiums far above the broad campaign to reach people ready to invest ten thousand dollars in a mutual fund.

168. Other smaller audiences that can be highly valuable include people who have already purchased from you, people who have something in a cart on your site, or people who meet certain criteria for your product—people who want to buy a new television, for example.

169. On the larger scale, larger audiences can be desirable and worth paying a high price for. In offline advertising, the classic example is the Super Bowl, which reaches most American households. In online advertising, a homepage takeover of a site like The Weather Channel, a top 10 worldwide site,⁶⁹ would be expensive but have tremendous reach to a broad audience.

170. In between these scales, we see the trade-offs that advertisers make in terms of size and audience quality. A tire company might purchase ads on the “road condition” portion of the site. A tire company that sells winter tires might purchase ads in this section during winter months north of the 40th parallel. If that company sold studded tires, they might further restrict their advertising to just states that permit studded tires. A promotion for a trade-in of a competitor’s tire might be targeted even more narrowly. At every step, there is a trade-off between the size of an audience and the amount of information known about the audience.

171. Advertisers use tools to help plan, manage, and execute their marketing campaigns. The advertising ad buying tools helps an advertiser execute the basics of digital marketing: uploading and storing ad creatives, decision-making for which ads to serve based on bidding and targeting, and tracking the performance of ads (like

⁶⁹ Semrush Blog. “Top 100: The Most Visited Websites in the US [2024 Top Websites Edition]” (April 2024). Accessed on June 4, 2024. <https://www.semrush.com/blog/most-visited-websites/>.

HIGHLY CONFIDENTIAL

clicks, impressions, conversions). Advertiser ad buying tools give advertisers control and visibility over their campaigns across multiple channels and publishers, allowing them to measure the effectiveness of their ads in real time.

172. On the demand side, advertisers often work with larger platforms to place ads at scale. As I have discussed above, today's marketers often place ads programmatically, which means that they also work with platforms that have built the infrastructure to purchase ad space inventory on a large scale.

173. The first demand side platforms (DSPs) emerged out of large advertising agencies who were buying large amounts of publisher inventory, on the order of billions of impressions per month, on behalf of many advertisers. Rather than having individual teams negotiating and purchasing inventory, sometimes resulting in teams negotiating against each other, the agencies formed the proto demand side platforms. Demand was aggregated at this level, allowing the entity to negotiate much better rates from the largest publishers.

174. DSPs are ad buying tools that allow advertisers to purchase ad inventory at a large scale in a system that is largely automated. There are three parts of buying ad inventory on DSPs. First, a marketer will develop a campaign and create advertising content, the ad creative. Next, the marketer will determine where the ad will be targeted and specify the spend or budget for the campaign. Third, the platform will search publisher inventory that meets the targeting criteria and can meet the budget. It then bids on the space for the ad, finalizes the bid, places the ad, and takes payment. Automation at so many stages of the marketing process reduces considerably the time advertisers take to place ads.

175. As the ad tech industry has evolved over time, many of the different tools that publishers and advertisers use to sell and serve ads have evolved, including ad servers on both the publisher and advertiser side, supply side platforms used by publishers, demand side platforms used by advertisers, ad buying tools, and even exchanges, which sometimes themselves bid in auctions, as though they are DSPs. Ultimately, it is most helpful to think of these tools as part of either buy side tools or sell side tools.

2. The Bifurcation of Tools for Large Advertisers and Small Advertisers

176. When it comes to tools for digital advertisers, the market has bifurcated into tools tailored specifically for large advertisers, such as DV360, and tools small advertisers, such as Google Ads. The customers for these tools have distinct needs and capabilities. Large advertisers typically operate at a scale that demands sophisticated tools capable of handling vast volumes of data and impressions. These advertisers often have dedicated internal teams and collaborate with external agencies to manage their marketing efforts. In fact, many large advertisers use different agencies

HIGHLY CONFIDENTIAL

for different advertising channels, taking advantage of the specialized skills developed by working deeply in one channel.⁷⁰ The tools they use, such as DSPs, provide advanced targeting, real-time bidding, and comprehensive analytics. These platforms allow for high levels of customization and control over campaigns, ensuring that large-scale efforts are optimized for maximum efficiency and effectiveness.

177. For small advertisers, marketing is often managed by individuals for whom it is not their primary job. These advertisers require tools that are straightforward, cost-effective, and easy to manage. Platforms like Google Ads⁷¹ and Facebook Ads⁷² have excelled in providing such solutions. Google Ads offers features like Smart Campaigns, which simplify the ad creation process and provide robust targeting options and analytics. Similarly, Facebook Ads Manager, including tools for Instagram, allows small advertisers to create targeted ad campaigns with a user-friendly interface and detailed performance metrics. These platforms balance automation with simplified control options, making them accessible to users with limited time and resources.

178. This distinction is by no means new. While large advertisers have always had complex marketing teams, small advertisers have migrated to digital from other channels, typically local advertising such as yellow pages, newspapers, and direct mail.⁷³ The key concept that allowed the small advertiser migration to digital is the monetization of the "long tail". Advertisers with niche products could never have profitably engaged in national advertising. With the advent of digital marketing, however, these advertisers can be connected to the right audience efficiently.⁷⁴

179. The differences in marketing tools reflect the distinct needs of large and small advertisers. Large advertisers, with their significant budgets and specialized teams, leverage advanced programmatic solutions to execute complex, multi-faceted campaigns. In contrast, small advertisers prioritize user-friendly platforms that offer essential targeting and analytics without requiring extensive resources. By

⁷⁰ Forbes. "16 Ways Companies Can Find Success With Multiple Agency Partners" (August 4, 2022). Accessed on June 4, 2024. <https://www.forbes.com/sites/forbesagencycouncil/2022/08/04/16-ways-companies-can-find-success-with-multiple-agency-partners/?sh=6a684bc97a81>; See also, Search Laboratory. "Eight steps to an integrated marketing strategy when you have multiple agencies" (February 17, 2020). Accessed on June 5, 2024. <https://www.searchlaboratory.com/us/2020/02/eight-steps-to-an-integrated-marketing-strategy-when-you-have-multiple-agencies/>.

⁷¹ Skiera, B., Bernd, Eckert, J., Jochen, & Hinz, O. Oliver (2010). "An analysis of the importance of the long tail in search engine marketing." Electronic Commerce Research and Applications vol., 9, no. 6. (2010), pgs. 488-494.

⁷² Facebook Ads is a tool for small advertisers advertising in the social media channel. As described in detail above, advertisers, including small advertisers, consider social a distinct channel from display because of the differences in audience, data for targeting, creative formats, and context within which the ads appear.

⁷³ Nowak, G., Cameron, G., Krugman, D. "How Local Advertisers Choose and Use advertising media," *Journal of Advertising Research* vol. 33, no. 6. 1993. pgs. 39-49.

⁷⁴ Denker, H., A. Enders, A., H. Hungenberg, H., P. Denker, and S. Mauch, S. "The Long Tail of Social Networking." *European Management Journal* vol. 26, no. 3. (2008). pgs. 199-211.

HIGHLY CONFIDENTIAL

understanding these differences, both types of advertisers can select the tools that best fit their goals, ensuring successful digital marketing strategies.

180. Marketers themselves also differ between these two groups. Large advertisers have dedicated internal teams focused on various aspects of marketing, including digital strategies, data analysis, and creative development. These teams often work alongside external agencies to amplify their efforts and bring additional expertise to their campaigns. On the other hand, small advertisers usually have individuals managing marketing as one of many responsibilities. This makes simplicity and ease of use crucial in the tools they select, as these users need to run effective campaigns without the depth of resources available to larger companies.⁷⁵

181. In summary, the bifurcation of the ad-buying tools market into tools for large and small advertisers reflects the varying needs and capacities of these groups. Large advertisers benefit from sophisticated, customizable platforms that handle large-scale operations, while small advertisers rely on accessible, user-friendly tools that facilitate efficient campaign management. This division ensures that advertisers of all sizes can effectively engage their target audiences and achieve their marketing objectives.

182. It is worth noting that this bifurcation is somewhat one directional. Small advertisers, unable to handle the cost or complexity of DV360, tend to use Google Ads exclusively. One might think that a symmetrical divide would exist, and that large advertisers would use DV360 exclusively to manage their display advertising. This is not the case, with [REDACTED] DV360 users also using Google Ads.⁷⁶ There are several reasons for this:

- Access to search: the overwhelming reason for the large, sophisticated advertisers who use DV360 to also use Google Ads is to participate in Google's search marketing. Advertisers who are managing their own search campaigns need to use Google Ads. And, while search is a distinct channel, once an advertiser is using Google Ads for search, the barrier is low to use it for parts of their display advertising.
- Enhanced targeting: Google Ads gives large advertisers access to targeting based on previous keyword searches, a powerful targeting opportunity akin to remarketing, discussed above.
- Resources: Large advertisers typically have many people working on marketing and thus have sufficient resources to manage both platforms.

⁷⁵ MIT Sloan Management Review. "From Niches to Riches: Anatomy of the Long Tail" (July 1, 2006). Accessed on June 3, 2024. <https://sloanreview.mit.edu/article/from-niches-to-riches-anatomy-of-the-long-tail/>.

⁷⁶ Figure 35, Expert Report of Robin S. Lee, Ph.D., December 22, 2023, GOOG-AT-MDL-C-000035792 at -952 (citing Google XP Data (DOJ RFP 7)).

HIGHLY CONFIDENTIAL

- Different access to inventory: in my experience, Google Ads provides much simpler access to YouTube than DV360.⁷⁷ Additionally, Google Ads is the primary way that most advertisers access the extensive Google Display Network.

E. Ad Exchanges

183. Ad exchanges are the place where advertisers and publishers come together to buy and sell ad inventory, in the form of impressions, in real-time auctions. Usually, to sell inventory in ad exchanges, publishers must meet a certain number of impressions. On advertising exchanges, advertisers place bids via DSPs and publishers make their inventory available for auction via SSPs.

184. While networks made it significantly easier for publishers to sell remnant inventory, the number of advertisers and publishers increased enough that by the mid-2000s, both advertisers and publishers, especially those buying and selling at scale, sought new technology to make serving ads even more streamlined.⁷⁸ Networks also had inefficiencies. There was a lack of price transparency, which translated to huge arbitrage opportunities for the networks. Furthermore, the lack of live bids left money on the table for publishers, while the uncertain delivery made it difficult for advertisers to make networks central to their buying strategies.

185. There are two main differences between ad networks and ad exchanges. The first is that ad networks are regarded as brokers for ads, while ad exchanges are open marketplaces, much like a stock exchange. The second difference is that ad exchanges are almost fully automated systems.

186. Right Media, founded in 2003, built the first ad exchange. It was revolutionary in that it allowed for real-time bidding, setting the foundation for what would become a significant shift in digital advertising.⁷⁹ Six years later, in 2009, Google launched its the DoubleClick Ad Exchange. Google created this exchange, now called AdX, by sending the inventory from its two massive networks, AdSense and the DoubleClick's network, into the exchange.

187. Brian O'Kelley, one of the people at the helm of building the first ad exchange, noted that the arrival of ad exchanges immediately meant that everyone was making larger profits from the buying and selling of programmatic ads. In his,

⁷⁷ DV360, on the other hand, provides access to certain features that are only available via DV360. See, Google Marketing Platform. "Take control of your campaigns" (undated). Accessed on June 6, 2024. <https://marketingplatform.google.com/about/display-video-360/>. ("Google Preferred, YouTube Reserve and TrueView inventory are all available for you to buy in Display & Video 360.").

⁷⁸ AdButler. "Ad Networks vs Ad Exchanges: The History of Programmatic Advertising" (March 15, 2021). Accessed on June 3, 2024. <https://www.adbutler.com/blog/article/ad-networks-vs-ad-exchanges-the-history-of-programmatic-advertising>.

⁷⁹ Tech Crunch. "How Mike Walrath Built Right Media and Sold it for \$850 Million" (April 12, 2011). Accessed on June 3, 2024. <https://techcrunch.com/2011/04/12/founder-stories-right-media-walrath/>.

HIGHLY CONFIDENTIAL

quoted in *New York Magazine: Intelligencer*, O’Kelley remarks that as soon as the ad exchange opened, “everyone’s yields went up at least 30 percent, instantly.” He goes on:

They made more money. We made more money. The advertisers made more money. The publishers made more money. Like, it was incredible. Everybody made more money. And it was just a perpetual motion machine.⁸⁰

188. OpenX began its journey as “PHP Ads” in 1998, later rebranding to “OpenAds” in 2007, and finally became “OpenX” in 2008.⁸¹ OpenX marked its position in the industry by transitioning from an ad network to an ad exchange, emphasizing transparency and real-time bidding capabilities. As ad exchanges evolved, other major platforms emerged, enhancing the ecosystem.

189. AppNexus (now part of Xandr, owned by Microsoft) quickly became known for its powerful programmatic ad technologies, providing tools for real-time sale and purchase of digital advertising.⁸²

190. The Rubicon Project (now merged with Telaria to form Magnite), which focused on automating the buying and selling of ads, grew significantly, supporting the trend towards more efficient and transparent marketplaces.⁸³

191. Google AdX, known formally as DoubleClick Ad Exchange, expanded under Google’s umbrella to become one of the most dominant ad exchanges globally, offering a vast range of inventory and advanced bidding features.⁸⁴

VIII. Google’s Position in the Digital Marketing Space

A. Opinions 7–10

192. As I have described above, the programmatic landscape of display advertising is shaped by intermediary ad tech platforms and tools, which streamline

⁸⁰ New York Magazine. “How to Succeed in Advertising (and Transform the Internet While You’re At It)” (undated). Accessed on June 3, 2024. <https://nymag.com/intelligencer/2018/05/right-media-creators-of-the-first-ad-exchange.html>.

⁸¹ AdButler. “Ad Networks vs Ad Exchanges: The History of Programmatic Advertising” (March 15, 2021). Accessed on June 3, 2024. <https://www.adbutler.com/blog/article/ad-networks-vs-ad-exchanges-the-history-of-programmatic-advertising>.

⁸² Adweek. “AT&T Unveils Xandr, Its Newly Rebranded Ad-Tech Unit” (September 25, 2018). Accessed on June 3, 2024. <https://www.adweek.com/programmatic/att-unveils-xandr-its-newly-rebranded-ad-tech-unit/>.

⁸³ AdExchanger. “Meet Magnite, The Post-Merger Name for Rubicon Project and Telaria” (June 30, 2020). Accessed on June 3, 2024. <https://www.adexchanger.com/platforms/meet-magnite-the-post-merger-name-for-rubicon-project-and-telaria/>.

⁸⁴ Google Blog. “A Year of the New DoubleClick Ad Exchange: Improving Large Publishers’ Returns” (January 16, 2011). Accessed on June 4, 2024. <https://googleblog.blogspot.com/2011/01/year-of-new-doubleclick-ad-exchange.html>.

HIGHLY CONFIDENTIAL

the transaction process between advertisers and publishers. These include publisher inventory management systems, publisher ad servers, advertising exchanges, advertiser ad servers, and advertiser buying tools. Over the years, the number of major ad tech competitors has decreased, and Google has become a dominant player in this space with products across all categories: DoubleClick for Publishers/Google Ad Manager, Google Ad Exchange, and DV 360. The consolidation of such tools within a single company raises concerns about conflicts of interest, particularly when the same entity operates on both the sell-side and buy-side, as well as within exchanges. Thus, I arrive at the following opinions:

193. **Opinion No. 7:** When faced with competitive threats, Google has strategically acquired competitors to maintain and enhance its market position. This approach has enabled Google to eliminate potential rivals and integrate valuable technologies, reinforcing its dominance in the ad tech ecosystem. Through these acquisitions, Google has built its dominant position in the display advertising market.

194. **Opinion No. 8:** Google provides and has provided platforms and tools in each of the foregoing categories, including its DoubleClick for Publishers (DFP) ad server, its Google Ad Exchange (AdX) exchange, its Google Ad Manager (GAM) ad server, and its DV 360 and Google Ads ad buying tools. Google is recognized in display advertising and the ad tech industry as the predominant player in publisher ad servers, ad exchanges, and ad-buying tools.

195. **Opinion No. 9:** There are inherent conflicts of interest when a single company provides both sell-side and buy-side platforms and tools, such as publisher ad servers and ad buying tools, respectively. The interests of publishers and advertisers are not generally aligned in a transaction for display advertising space. Conflicts of interest can harm the transparency and fairness of the auction process. These conflicts of interest can be a problem when participants who do not have good alternatives.

196. **Opinion No. 10:** Additional conflicts of interest arise when the same company is also in the exchange business. Google is such a company, and the digital advertising and ad tech industries generally recognize the existence of Google's multiple conflicts of interest.

B. Introduction

197. We have had a chance to look at the emergence of the digital marketing and ad tech ecosystems, now we'll turn in this section to an overview of how Google positioned itself in the ad tech ecosystem. What we see, in general, is that over time, Google acquired ad tech companies strategically. Each of its acquisitions brought unique capabilities, technologies, and talent into Google's fold, and propelled its unique, and dominant, position in the advertising ecosystem. Its acquisitions also gave it a dominant informational position so that it could view the ecosystem from

HIGHLY CONFIDENTIAL

multiple perspectives at once and make strategic buying and selling positions accordingly.

C. A Brief History of Google's Acquisitions and Rise to Power in the Ad Tech Industry


198. Google's journey in developing its search marketing business began in the late 1990s, following the company's founding by Larry Page and Sergey Brin in 1998. The core of Google's success lay in its PageRank algorithm, which provided more relevant search results than its competitors. This innovation attracted a growing user base, laying the foundation for Google to monetize its search engine.

199. In 2000, Google launched AdWords, a self-service ad program that allowed businesses to display ads alongside search results. Initially, AdWords operated on a cost-per-impression (CPM) model, but it soon transitioned to a cost-per-click (CPC) model, which proved to be more profitable and to allow Google to maintain greater control over their ecosystem. Advertisers only paid when users clicked on their ads, ensuring that their marketing budgets were spent on interested audiences. The introduction of the Quality Score system, which evaluated ad relevance and landing page quality, further optimized the user experience and ad performance. This approach, combined with Google's vast user data, enabled highly targeted advertising and allowed Google to grow to dominate the search advertising business. In 2002, Google had 29% of search traffic⁸⁵ and by 2005 Google was number 1, with 37% of traffic.⁸⁶ The rise was meteoric and by 2015 Google controlled around 90% of search share.⁸⁷

⁸⁵ Manning Search Marketing. "Top Search Engines from 2002 to 2005" (February 6, 2014). Accessed on June 3, 2024. <https://www.manningmarketing.com/articles/top-search-engines-2002-2005/>.

⁸⁶ Comscore. "Comscore Reports July 2005 Search Engine Rankings" (August 19, 2005). Accessed on June 3, 2024. <https://ir.comscore.com/news-releases/news-release-details/comscore-reports-july-2005-search-engine-rankings>.

⁸⁷ StatCounter. "Desktop Search Engine Market Share Worldwide" (undated). Accessed on June 3, 2024. <https://gs.statcounter.com/search-engine-statistics/216573/worldwide-market-share/desktop/worldwide/#monthly-201501-202301-of-search-engines/>.

HIGHLY CONFIDENTIAL

200. Google's acquisitions in the ad tech space are extensive and reflect Google's approach to gaining control over specific tools and forms of information that could strengthen its overall position in the ad tech ecosystem. From 2003-2023, Google acquired 255 companies, at least fourteen of which are in the ad tech landscape.⁸⁸ I do not survey all acquisitions and will instead consider some key acquisitions that are relevant to its role in the advertising space in general, and the display marketing space in particular.

⁸⁸ CrunchBase. (November 28, 2022). "Price of selected acquisitions by Google as of November 2022 (in million U.S. dollars)." Graph. June 07, 2024. Statista. Accessed June 07, 2024. <https://www.statista.com/statistics/192300/price-of-selected-acquisitions-by-google/>; See also, Tracxn. "Acquisitions by Google" (April 8, 2024). Accessed on June 7, 2024. [https://tracxn.com/d/acquisitions/acquisitions-by-google/_8zKOTBXR934x_3BSnEVSrmsu3RZtEv3AorQtuBb2Yk#:~:text=Acquisitions%20by%20Google,-1998%E2%80%A2Mountain&text=Google%20has%20made%20260%20acquisitions,Consumer%20Digital%20%20US%20and%20others](https://tracxn.com/d/acquisitions/acquisitions-by-google/_8zKOTBXR934x_3BSnEVSrmsu3RZtEv3AorQtuBb2Yk#:~:text=Acquisitions%20by%20Google,-1998%E2%80%A2Mountain&text=Google%20has%20made%20260%20acquisitions,Consumer%20Digital%20%20US%20and%20others;); Sales Tech Series. "Google's Acquisition Spree: A Journey Through 2023's Deals" (December 20, 2023). Accessed on June 7, 2024. <https://salestechstar.com/sales-engagement/googles-acquisition-spree-a-journey-through-2023s-deals/#:~:text=Having%20made%20an%20incredible%2056,fueled%20its%20expansion%20and%20influence>

HIGHLY CONFIDENTIAL

201. Google's foray into display advertising began with the acquisition of Applied Semantics in April 2003.⁸⁹ Based in California, Applied Semantics specialized in contextual advertising technology, enabling ads to be targeted based on the content of webpages. While the Applied Semantics technology was undoubtedly used for Google's search product, Ad Words, this technology was the fundamental underpinning of AdSense, Google's first display advertising product.⁹⁰ As adweek.com observed at the time, this acquisition shows Google "continuing its efforts to push advertising beyond just search engine paid placement and onto actual web pages."⁹¹ This acquisition laid the foundation for Google's contextual advertising platform, AdSense, empowering advertisers to reach their target audience with precision.⁹²

202. Also in 2003, Google acquired Sprinks online ad service from Primedia, giving it a larger share of the emerging contextual advertising space, which is advertising that focuses on the context or place in which advertisements appear, as opposed to targeted advertising, which focuses on reaching users with certain attributes. As MediaPost.com reported at the time, this deal "made Google the exclusive provider of contextual ad sales—also known as pay-per-click advertising sales—for Primedia's About.com unit and for most of Primedia's Consumer Media and Magazines Group websites."⁹³ At this point, Google had not yet gone public, but as MediaPost also

⁸⁹ Adweek. "Google Acquires Applied Semantics" (April 24, 2003). Accessed on June 3, 2024. <https://www.adweek.com/brand-marketing/google-acquires-applied-semantics-63588/>. See also, CBInsights. "The Google Acquisition Tracker" (undated). Accessed on June 5, 2024. <https://www.cbinsights.com/research/google-acquisitions>.

⁹⁰ As mentioned in Section V, digital marketing channels are characterized by their unique contexts and mechanisms of engagement, extending beyond the format of the ads they display. As I mentioned above, we see this distinction particularly clearly when comparing channels like social media marketing and display advertising. Despite sometimes showcasing similar creatives, these channels operate within fundamentally different environments and user experiences. Social media marketing leverages the interactive, community-oriented nature of social platforms to foster engagement and brand loyalty, while display advertising typically targets users through static or dynamic ads placed across a network of websites. Each channel has its own strategies, user expectations, and measurement metrics, which contribute to its distinct identity within the digital marketing ecosystem. Similarly, Google's AdSense and AdWords, though both initially focused on text ads, represent different facets of digital advertising. AdWords, now known as Google Ads, primarily targets users through search engine results pages. Advertisers bid on keywords relevant to their products or services, and their ads appear alongside organic search results when users enter those keywords. In contrast, AdSense is centered around contextual advertising on a diverse network of third-party websites. Website owners enroll in the AdSense program to display ads that are relevant to their site's content, thereby generating revenue based on user interactions with those ads. Unlike AdWords, which targets users based on their search queries, AdSense targets users based on the content they are consuming on various web pages.

⁹¹ Adweek. "Google Acquires Applied Semantics" (April 24, 2003). Accessed on June 3, 2024. <https://www.adweek.com/brand-marketing/google-acquires-applied-semantics-63588/>.

⁹² AllThingsD. "Ten Years Later: Lessons From the Applied Semantics' Google Acquisition" (April 22, 2013). Accessed on June 3, 2024. <https://allthingsd.com/20130422/ten-years-later-lessons-from-the-applied-semantics-google-acquisition/>.

⁹³ Online Media Daily. "Google Acquires Sprinks, Consolidates Contextual Ad Market" (October 26, 2003). Accessed on June 3, 2024. <https://www.mediapost.com/publications/article/18214/google-acquires-sprinks-consolidates-contextual-a.html>.

HIGHLY CONFIDENTIAL

observed, many ad tech professionals began to speculate “that Google would soon mount an initial public offering that would likely raise capital to expand its base even further.”⁹⁴ Google ultimately went public on August 19, 2004, netting approximately \$1.67 billion.⁹⁵

203. In late 2005, Google made a strategic move to expand its advertising reach beyond the digital realm and acquired dMarc Broadcasting. Based in California, dMarc specialized in radio advertising automation, enabling advertisers to target audiences through radio channels.⁹⁶ As informationweek.com wrote in 2006, “Google said it plans to create a new radio ad distribution channel by integrating dMarc’s technology into Google’s AdWords platform. dMarc assists broadcasters by automatically scheduling and placing advertising.”⁹⁷ This acquisition diversified Google’s advertising portfolio, offering advertisers a multichannel approach to reaching their target demographics. Google’s acquisitions suggest in general that it always sought different channels in the ad tech market because of their distinct and diverse revenue streams.

204. In October 2006, Google acquired YouTube.⁹⁸ While YouTube functions in digital marketing as primarily a publisher of video content, this acquisition was notable for the dominant position it gave Google in a novel marketing channel. By 2023 YouTube, with an annual revenue of \$31.5 billion, accounted for 10% of Google’s total revenue.

205. In 2007, Google extended its reach into the realm of in-game advertising with the acquisition of Adscape Media.⁹⁹ Based in California, Adscape Media pioneered dynamic advertising within video games, opening up new avenues for advertisers to engage with consumers in immersive virtual worlds.

206. In 2008, Google made one of its most significant acquisitions, DoubleClick, for \$3.1 billion in cash. As explained above, DoubleClick was a pioneer and a powerhouse in online advertising technology. Google gained access to DoubleClick’s advertising software and its extensive relationships with web publishers, advertisers,

⁹⁴ Online Media Daily. “Google Acquires Sprinks, Consolidates Contextual Ad Market” (October 26, 2003). Accessed on June 3, 2024. <https://www.mediapost.com/publications/article/18214/google-acquires-sprinks-consolidates-contextual-a.html>.

⁹⁵ CNN Money. “Google goes low: IPO set at \$85 a share, low end of revised range, with trading set to start Thursday” (August 19, 2004). Accessed on June 6, 2024. <https://money.cnn.com/2004/08/18/technology/googleipo/>.

⁹⁶ Information Week. “Google Acquires dMarc Radio for \$102 Million” (January 17, 2006). Accessed on June 3, 2024. <https://www.informationweek.com/it-leadership/google-acquires-dmarc-radio-for-102-million#close-modal>.

⁹⁷ *Id.*

⁹⁸ NBC News. Google Buys YouTube for \$1.65 Billion” (October 9 2006). Accessed on June 3, 2024. <https://www.nbcnews.com/id/wbna15196982>.

⁹⁹ Google. “Let the passion continue! We’re acquiring Adscape” (March 16, 2007). Accessed on June 6, 2024. <https://googleblog.blogspot.com/2007/03/let-passion-continue-were-acquiring.html>

HIGHLY CONFIDENTIAL

and advertising agencies.¹⁰⁰ As the New York Times reported in 2007, “[t]he sale offers Google access to DoubleClick’s advertisement software and, more importantly, its relationships with Web publishers, advertisers and advertising agencies.” It also significantly expanded Google’s foothold in online advertising into display ads, the area where DoubleClick is strongest.”¹⁰¹ When Google purchased DoubleClick, the advertiser ad server had a 63% share in the United States.¹⁰²

103

207. With its suite of ad management and serving solutions, DoubleClick bolstered Google’s ad infrastructure, enabling advertisers to manage and optimize their campaigns with unprecedented efficiency. This acquisition solidified Google’s position as a dominant force in display advertising. As the New York Times also reported during that period, the acquisition occurred in competition with Microsoft, and had it not gone through, it would have been a blow to Google.

208. The New York Times also observed Google’s increasing reach in display advertising and other channels: “Google has been expanding its AdSense network into video and display ads online and is selling ads to a limited degree on television, newspapers and radio.” The New York Times continues, writing that the acquisition, “...raises questions about how Google will manage its existing business and that of the new DoubleClick unit while avoiding conflicts of interest. If DoubleClick’s existing clients start to feel that Google is using DoubleClick’s relationships to further its own ad network, some Web publishers or advertisers might jump ship.”¹⁰⁴

209. At least as early as 2007, then, Google’s acquisitions were setting it up as uniquely powerful across the ad tech ecosystem. During my time at Atlas, I participated in sales calls where we raised concerns about using DoubleClick, arguing that it would give Google too much control over the advertiser’s marketing.

210. Perhaps most significantly, DoubleClick’s introduction of a Nasdaq-like exchange for online ads presented a lucrative opportunity for Google to expand its

¹⁰⁰ New York Times. “Google Buys DoubleClick for \$3.1 Billion” (April 14, 2007). Accessed on June 3, 2024. <https://www.nytimes.com/2007/04/14/technology/14DoubleClick.html>.

¹⁰¹ New York Times. “Google Buys DoubleClick for \$3.1 Billion” (April 14, 2007). Accessed on June 3, 2024. <https://www.nytimes.com/2007/04/14/technology/14DoubleClick.html>.

¹⁰² AEI-Brookings Joint Center for Regulatory Studies. “An Antitrust of Google’s Proposed Acquisition of DoubleClick” (September 2007). Accessed on June 4, 2024. https://www.brookings.edu/wp-content/uploads/2016/06/09useconomics_hahn.pdf.

¹⁰³ SimilarTech. “Publisher Ad Server” (undated). Accessed on June 6, 2024. <https://www.similartech.com/categories/publisher-ad-server>.

¹⁰⁴ The New York Times. “Google Buys DoubleClick for \$3.1 Billion” (April 14, 2007). Accessed on June 4, 2024. (<https://www.nytimes.com/2007/04/14/technology/14DoubleClick.html>; See also Federal Trade Commission. “Dissenting Statement of Commissioner Harbour In the Matter of Google/DoubleClick” (December 20, 2007). Accessed on June 5, 2024. <https://www.ftc.gov/legal-library/browse/cases-proceedings/public-statements/dissenting-statement-commissioner-harbour-matter-googledoubleclick>).

HIGHLY CONFIDENTIAL

presence in the display advertising business. As the New York Times reported, “DoubleClick’s chief executive, David Rosenblatt, said a few weeks ago that a new system it had developed for the buying and selling of online ads would probably become the chief moneymaker within five years.”

211. Many people remarked on the acquisition at the time, not the least being FTC Commissioner Pamela Harbor who objected to the deal, and who warned in 2007 that the merger would lessen competition and consumer protections related to privacy.¹⁰⁵

212. In recent deposition testimony for the current Department of Justice civil antitrust suit against Google for monopolizing search and search advertising [REDACTED] was asked about the impact of Google’s acquisition of DoubleClick on competition in the ad tech industry. [REDACTED]

213. After the acquisition, [REDACTED]

106

214. [REDACTED] on the way that the DoubleClick acquisition impacted competition, stating [REDACTED]

107

215. After its acquisition of DoubleClick, Google’s continued acquisitions repeat a familiar pattern. In 2009, as mobile usage skyrocketed, Google recognized the importance of mobile advertising and swiftly acquired AdMob, a leading mobile ad platform based in California, for \$750 million. With AdMob’s technology and reach, Google gained a foothold in the burgeoning mobile advertising market, allowing advertisers to target users on smartphones and tablets effectively. As zdnet.com wrote at the time, “with the move, Google is making a big bet on mobile display advertising. Google, which already occupies the market for mobile text ads, appears to be

¹⁰⁵ Tech Policy Press. “How Three Mergers Buttressed Google’s Ad Tech Monopoly, Per DOJ” (March 9, 2023). Accessed on June 4, 2024. “<https://www.techpolicy.press/how-three-mergers-buttressed-googles-ad-tech-monopoly-per-doj/>.

¹⁰⁶ Deposition of [REDACTED] 72:6-74:15. [REDACTED]

¹⁰⁷ [REDACTED] 74:16-75:21.

HIGHLY CONFIDENTIAL

purchasing AdMob in an acquisition akin to its DoubleClick purchase. The message: Google wants to be a big display ad player on the Internet and mobile.”¹⁰⁸

216. Adnet.com’s remarks on the acquisition of AdMob echoes some of the reasoning in the New York Times report regarding DoubleClick. As adnet.com writes at the time to conclude its report on the AdMob acquisition:

It doesn’t take a rocket scientist to figure out that Google is positioning itself in the same center of the ad universe. AdMob gets a nice payday, more resources and a big partner. Given the alternative—eventually being run over by Google—it’s hard to turn down the deal.¹⁰⁹

217. In sum, as early as sixteen years ago, the sense in the industry was that ad tech companies face two choices: be bought by Google or be destroyed by it.

218. In late 2009, Google’s quest for personalized advertising led to the acquisition of Teracent Corporation. Based in California, Teracent specialized in dynamic ad optimization, leveraging machine learning algorithms to deliver personalized display ads tailored to individual user’s preferences and behaviors. As Google stated with the announcement of the acquisition, “We think that this technology has great potential to improve display advertising on the web.”¹¹⁰ Google also elaborated on the specific capabilities that Teracent would give them for ad targeting:

Teracent’s technology can pick and choose from literally thousands of creative elements of a display ad in real-time — tweaking images, products, messages or colors. These elements can be optimized depending on factors like geographic location, language, the content of the website, the time of day or the past performance of different ads. This technology can help advertisers get better results from their display ad campaigns. In turn, this enables publishers to make more money from their ad space and delivers web users better ads and more ad-funded web content.¹¹¹

219. In a postscript, an industry commentator notes this acquisition “put Google on par with Yahoo in terms of its ability to deliver dynamic display advertising on the PC and do a range of targeting types, including behavioral targeting.” The commentator goes on to note that the mantra in the mobile world has been “right ad,

¹⁰⁸ ZDNET. “Google Picks Up AdMob for \$750 million; Target Mobile Display Ads” (November 9, 2009). Accessed on June 4, 2024. <https://www.zdnet.com/article/google-picks-up-admob-for-750-million-targets-mobile-display-ads/>.

¹⁰⁹ ZDNET. “Google Picks Up AdMob for \$750 million; Target Mobile Display Ads” (November 9, 2009). Accessed on June 4, 2024. <https://www.zdnet.com/article/google-picks-up-admob-for-750-million-targets-mobile-display-ads/>.

¹¹⁰ Google Official Blog. “Displaying the best display ad with Teracent” (November 23, 2009). Accessed on June 6, 2024. <https://googleblog.blogspot.com/2009/11/displaying-best-display-ad-with.html>.

¹¹¹ *Id.*

HIGHLY CONFIDENTIAL

right time, right place,” and “[t]hat’s not possible without this dynamic capability, which Yahoo has in mobile (in fact provided by Teracent, though perhaps not for much longer).” He concludes, “think about the combination of this technology and AdMob and you start to see how Google will be able to deliver something pretty compelling in the mobile display arena — an ad platform that begins to truly leverage the additional capabilities and location-awareness of mobile.” Again, we can see, even this early in the game, industry commentators expecting that Google will leverage the combination of technologies it is acquiring, and that it will become a display advertising powerhouse.

220. In March 2010, Google acquired Episodic.¹¹² Based in California, Episodic was a pioneering force in interactive video technology. Its innovative solutions enabled seamless integration of interactive elements into video content, enhancing user engagement and interactivity. By acquiring Episodic, Google fortified its position in the realm of video advertising, leveraging its technology to deliver immersive and interactive ad experiences. This acquisition empowered Google to offer advertisers innovative ways to captivate audiences and drive engagement, thereby strengthening its competitive edge in the advertising landscape.

221. In June 2010, Google expanded further its advertising prowess with the acquisition of Invite Media, a demand-side platform based in New York City. Invite Media, founded in 2007, was among the pioneering players in the DSP space.¹¹³ The acquisition was a strategic move by Google to bolster its display ad buying capabilities. The deal had been in the works since late 2009, but, at the time, industry analysts remarked that Google delayed finalizing it to avoid regulatory scrutiny from the FTC amid its AdMob acquisition.¹¹⁴ By integrating Invite Media's tools into its stack, including DoubleClick and the Google Ad Exchange, Google aimed to offer advertisers enhanced targeting options and real-time bidding capabilities. However, concerns also arose about impartiality, once again, with competitors raising questions about potential bias towards Google's own inventory.

222. In June 2011, Google bolstered its publisher-side offerings with the acquisition of AdMeld. AdMeld, founded in 2007, and headquartered in New York, operated one of the largest SSPs in the display advertising industry. AdMeld provided publishers with tools to maximize their ad revenue through advanced yield optimization and inventory management solutions. It catered to large, premium publishers

¹¹² Tech Crunch. “Google Acquires Online Video Hosting Platform Episodic” (April 2, 2010). Accessed on June 4, 2024. <https://techcrunch.com/2010/04/02/google-acquires-online-video-hosting-platform-episodic/>

¹¹³ ClickZ. “Google Acquires Invite Media, First Buy-Out of a DSP” (June 2, 2010). Accessed on June 4, 2024. <https://www.clickz.com/google-acquires-invite-media-first-buy-out-of-a-dsp/47497/>.

¹¹⁴ ClickZ. “Google Acquires Invite Media, First Buy-Out of a DSP” (June 2, 2010). Accessed on June 4, 2024. <https://www.clickz.com/google-acquires-invite-media-first-buy-out-of-a-dsp/47497/>.

HIGHLY CONFIDENTIAL

with over 50 million remnant impressions per month.¹¹⁵ Its clients included the Huffington Post and WorldNow, for instance.¹¹⁶ The acquisition underwent rigorous scrutiny by the Department of Justice's Antitrust Division.¹¹⁷ The DOJ ultimately allowed the AdMeld acquisition, but it stated emphatically that it would continue to enforce all antitrust activity. With AdMeld onboard, Google strengthened its relationships with large, premium publishers, further cementing its position in the digital advertising ecosystem.

223. The acquisition of AdMeld marked an important milestone for Google in its display advertising journey. I elaborate on this milestone further down, but this acquisition marks the first time that Google held a dominant position in the three critical technologies that enable digital marketing: advertiser tools, publisher tools, and ad exchanges.

224. In July 2012, Google continued its expansion in the creative realm when it acquired California-based Cuban Council. Cuban Council was a creative agency renowned for its digital design and advertising expertise. While specific details of the acquisition were not disclosed, Google's move to acquire Cuban Council signaled its commitment to enhancing its creative capabilities in advertising.¹¹⁸

225. In July 2012, recognizing the growing influence of social media in advertising, Google acquired WildFire, based in California. WildFire's social marketing platform empowered advertisers to engage with audiences across various social channels, enabling them to amplify their brand presence and drive conversions.¹¹⁹ This acquisition enriched Google's advertising offerings, integrating social media into its comprehensive suite of marketing solutions.

226. In November 2012, Google acquired Incentive Targeting, a Massachusetts company that leveraged data analytics to deliver targeted promotions and offers to consumers, driving sales and customer loyalty. Google's acquisition of Incentive

¹¹⁵ AdExchanger. "Optimizing AdMeld Style With Co-Founder Ben Barokas" (April 8, 2009). Accessed on June 4, 2024. <https://www.adexchanger.com/yield-management-tools/admeld-ad-exchange-optimization>.

¹¹⁶ *Id.*

¹¹⁷ Office of Public Affairs. Statement of the Department of Justice's Antitrust Division on Its Decision to Close Its Investigation of Google Inc.'s Acquisition of Admeli Inc." (December 2, 2011). Accessed on June 4, 2024. <https://www.justice.gov/opa/pr/statement-department-justices-antitrust-division-its-decision-close-its-investigation-google>.

¹¹⁸ The Verge. "Google Hires Talent from Cuban Council Firm to Bolster Google+ Design" (July 18, 2012). Accessed on June 4, 2024. <https://www.theverge.com/2012/7/18/3166915/google-cuban-council-talent-acquisition>.

¹¹⁹ Search Engine Journal. "Google Spends \$400 Million to Purchase Social Marketing Startup Wildfire" (August 2, 2012). Accessed on June 4, 2024. <https://www.searchenginejournal.com/google-acquisition-wildfire-social/46958/>.

HIGHLY CONFIDENTIAL

Targeting bolstered its arsenal of advertising tools, particularly in the realm of local and retail advertising, extending its foothold in the retail advertising market.¹²⁰

227. In February 2015, Google acquired Red Hot Labs, a California-based mobile gaming startup with an unconventional approach to mobile advertising and user acquisition. Specializing in user acquisition and growth hacking strategies, Red Hot Labs offered insights and expertise in mobile advertising optimization.¹²¹ Google's acquisition of Red Hot Labs signaled its strategic focus on the mobile advertising market and its commitment to enhancing its mobile ad offerings. This acquisition positioned Google as a leader in the rapidly evolving mobile advertising landscape, driving innovation and growth in the mobile ad market.

228. In October 2016, Google acquired FameBit, a leading influencer marketing platform that facilitated collaborations between brands and digital creators on YouTube and other social media platforms. In Google's ongoing quest to redefine advertising paradigms, the acquisition of FameBit marked a strategic move into the realm of influencer marketing.¹²²

1. Google's Ad Tech Tools and Platforms

229. Google's acquisitions over the past two decades reveal a strategic approach to buy tools in every facet of the ad tech ecosystem. By targeting key companies and technologies, Google has ensured its presence and influence across the entire digital advertising landscape.

230. The DoubleClick acquisition created Google's position in advertiser ad serving, publisher ad serving, and ad exchange technology. Before acquiring DoubleClick, Google was primarily a search publisher, with a side business in display via AdSense. The DoubleClick acquisition was a watershed moment in solidifying its control over publisher tools and SSPs. DoubleClick's ad management and serving solutions allowed Google to offer comprehensive ad delivery and optimization services to publishers. Additionally, the acquisition of AdMeld in 2011 allowed Google to offer advanced yield optimization and inventory management solutions. These acquisitions provided Google with unequaled capabilities to manage and optimize ad inventory from the publisher's perspective, making it a central player in the supply side of digital advertising.

¹²⁰ Tech Crunch. "Confirmed: Google Acquires Incentive Targeting to Power Super Targeted, Personalized Coupon Programs" (November 28, 2012). Accessed on June 4, 2024.

<https://techcrunch.com/2012/11/28/google-acquires-incentive-targeting-to-power-targeted-coupon-programs/>.

¹²¹ TechCrunch. "Google Acquires Facebook Marketing Startup Toro" (February 24, 2015). Accessed on June 6, 2024. <https://techcrunch.com/2015/02/24/google-acquires-toro/>.

¹²² TechCrunch. "Google acquires FameBit to connect YouTube creators with marketers" (October 11, 2016). Accessed on June 6, 2024. <https://techcrunch.com/2016/10/11/google-acquires-famebit-to-connect-youtube-creators-with-brands/>.

HIGHLY CONFIDENTIAL

231. On the advertiser side, the DoubleClick acquisition combined with the acquisition of Invite Media marked its entry into the DSP market. Invite Media's real-time bidding and advanced targeting capabilities enabled Google to offer advertisers sophisticated tools to purchase ad inventory at scale. The acquisition of AdMob further expanded Google's reach into in-app advertising, allowing advertisers to target users on mobile devices effectively.

232. Google's acquisition of DoubleClick also brought it into the ad exchange arena. DoubleClick's ad exchange capabilities allowed Google to facilitate real-time bidding and transactions between advertisers and publishers. This integration enabled Google to create a seamless and efficient marketplace for digital ad inventory, further cementing its role as a key intermediary in the ad tech ecosystem. The additional acquisitions of companies like AdMob and AdMeld further integrated these capabilities, allowing Google to streamline the buying and selling process across multiple channels.

233. This trifecta continued to ring alarm bells regarding Google's position in the ad tech market. As noted above, along with other industry observers, the FTC did not approve the DoubleClick acquisition by a unanimous vote. As techpolicy.press observed at the time, "Three deals — the purchases of DoubleClick, Invite Media, and AdMeld — 'set the stage' for the tech giant to 'control and manipulate' the selling and buying of digital ads, says DOJ." As someone deeply enmeshed in this industry for twenty-five years, it is easy to see this present lawsuit as a case of "chickens coming home to roost." Google has not successfully entered any marketing channels other than search via home-grown products. If the DoubleClick acquisition had not been approved, this present lawsuit would not be necessary.

234. In light of my remarks in Sections IV, V, and VI, we can see that there are three primary ad tech products that work together to enable digital marketing: publisher inventory management systems' ad servers, advertiser ad servers and buying tools, and ad exchanges. As techpolicy.press commentators note, Google's acquisitions represent Google's primary competition in each of these areas (DoubleClick for exchanges and advertiser ad serving, Invite Media for ad buying tools, and AdMeld for publisher ad servers). As techpolicy.press also notes¹²³:

[A]fter merging all these companies into Google's own ad tech stack, the tech giant owns the ad exchange AdX, the publisher ad server Google Ad Manager, and two ad buying tools, DV360 (which serves large advertisers) and Google Ads (for smaller advertisers). In each of these product segments, Google's market share measured by revenues ranges from 40 to 90 percent, DOJ revealed in the suit.

¹²³ Tech Policy. Press. "How Three Mergers Buttressed Google's Ad Tech Monopoly, Per DOJ" (March 9, 2023). Accessed on June 4, 2024. "<https://www.techpolicy.press/how-three-mergers-buttressed-googles-ad-tech-monopoly-per-doj/>".

HIGHLY CONFIDENTIAL

235. We can see Google positioning itself here as a dominant player who has access to all sides of the buying and selling process, as well as owning the auction house itself. With these acquisitions in place, Google was positioned to expand its position in the ad tech market significantly and in a singular way.

236. Google's dominance in the ad tech marketplace is also confirmed by a number of other experts who have testified on this matter. [REDACTED]
[REDACTED] was asked.
[REDACTED]
[REDACTED]

237. In his deposition, [REDACTED]
[REDACTED]

125

238. In his deposition, [REDACTED], remarked [REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

2. DFP (GAM)—Publisher Ad Server

239. Currently, Google's Publisher Ad Server is called Google Ad Manager (GAM). It was renamed in 2018, though many of its features remain the same. Many in the industry refer to GAM as DFP, since DFP was the earlier product and more widely known.

3. AdX (GAM)—Exchange

240. The integration of DoubleClick's ad exchange capabilities marked Google's entry into the ad exchange arena, facilitating real-time bidding and transactions between advertisers and publishers. This acquisition enabled Google to create a

¹²⁴ Deposition of [REDACTED] 19:1-19:12. [REDACTED].

¹²⁵ Deposition of [REDACTED] 83:8-83:11. [REDACTED].

¹²⁶ Deposition of [REDACTED] 73:3-73:5. [REDACTED]

HIGHLY CONFIDENTIAL

unified marketplace for digital ad inventory, further strengthening its role as a key intermediary in the ad tech ecosystem. AdX, which requires GAM or a partner with AdX access¹²⁷, controls a substantial portion of programmatic inventory.

4. Google Ads—Ad Buying Tool for Small Advertisers

241. Google Ads, initially known as AdWords, was developed to cater to small advertisers, providing a straightforward platform with basic targeting options and analytics. Google Ads provides access to advertising on Google Search as well, so many large advertisers use Google Ads in a limited capacity to take advantage of this feature, though most large advertisers use DV360

5. DV360—Ad Buying Tool for Large Advertisers

242. The acquisition of Invite Media, combined with the DoubleClick for Advertisers (DFA) product from the DoubleClick acquisition, allowed Google to introduce DV360, a DSP and ad buying tool designed for large advertisers.

6. Google's Other Digital Ad Tools and Platforms

243. Beyond DoubleClick, Invite Media, and AdMeld, Google has acquired several other key companies to enhance its ad tech capabilities. AdMob, for instance, extended Google's reach into mobile advertising, allowing advertisers to effectively target users on mobile devices. Google Analytics, created from the acquisition of Urchin Software in 2005,¹²⁸ provides detailed insights into website traffic and user behavior, aiding advertisers in optimizing their campaigns while providing Google with a nearly complete view of activity on the advertiser's website. Additionally, Google Tag Manager simplifies the process of managing and deploying marketing tags on websites without requiring extensive coding, while again giving Google unrivaled data access. These tools, among others, integrate to form a comprehensive suite that supports various aspects of digital advertising and data analysis.

D. The Potential Conflicts of Interest and Self-Preferencing Opportunities Inherent in Google's Position

244. Google's extensive acquisitions in the ad tech space have positioned it as a dominant player with control over all facets of display advertising. This consolidation of power presents potential conflicts of interest, particularly when a single entity oversees both the buy-side and sell-side platforms, as well as the exchange where transactions occur. This dual role leads to self-preferencing behaviors, where Google's platforms might favor its own services and inventory over competitors', thereby disadvantaging other players in the market.

¹²⁷ Google Certified Publishing Partner. "Find the perfect partner." (undated). Accessed on June 6, 2024. <https://www.google.com/ads/publisher/partners/find-a-partner/>.

¹²⁸ New York Times. "Google Acquires Urchin Software" (March 29, 2005). Accessed on June 4, 2024. <https://www.nytimes.com/2005/03/29/technology/google-acquires-urchin-software.html>.

HIGHLY CONFIDENTIAL

245. The integration of Google's ad tech tools—such as Google Ads for small advertisers, DV360 for large advertisers, and Google Ad Manager for publishers—demonstrates how intertwined its operations have become. By owning the auction house through AdX and controlling significant ad server technologies via DoubleClick, Google has unparalleled access to data and can influence market dynamics in ways that benefit itself at the expense of the advertisers and publishers its tools purport to serve. This consolidation raises significant concerns, as it may stifle competition, limit innovation, and create an uneven playing field for other participants in the digital advertising market.

246. Moreover, the asymmetry of information that Google possesses due to its comprehensive data collection capabilities further exacerbates these concerns. Advertisers and publishers depend on the transparency and fairness of the auction process to optimize their strategies and maximize revenues. However, when the entity controlling critical parts of this ecosystem also participates in the bidding process, it creates an inherent conflict of interest that can undermine the integrity of the market.

IX. The Importance of Information and Transparency in the Programmatic Display

A. Opinions 11–16

247. The transparency and overall fairness of the programmatic display auction process hinge on the availability and accessibility of critical information. Advertisers require comprehensive data about the auction mechanics, algorithms, audience characteristics, ad placement specifics, intermediary fees, and ad performance to optimize their strategies and platform choices. Similarly, publishers need detailed insights into these areas to maximize their revenues and set effective price floors. Any restriction, unequal access, or misrepresentation of this essential information can significantly impact the fairness and efficiency of the auction process, undermining trust and competitiveness in the market. As such, my opinions from this section are the following:

248. **Opinion No. 11:** Advertisers and publishers depend on transparency and fairness when they engage in the programmatic website display auction process. These in turn depend in large part upon the nature and extent of the available information regarding that auction and the degree of the participants' and intermediaries' access to necessary information.

249. **Opinion No. 12:** To maximize the cost-effectiveness of their purchase of programmatic display advertising, and to optimize their auction-related strategies and platform choices, advertisers typically need: (a) data and information regarding the mechanics and rules of the auction process; (b) an understanding of the algorithms employed by the intermediary platforms, tools, and exchanges; (c) data and

HIGHLY CONFIDENTIAL

information about the website visitors who will ultimately receive the display ad; (d) information about the space (i.e., impression) where the advertisement will be displayed; (e) information about the commission, share, take rate, mark-up, or other portion of their payment that is retained by intermediaries; and (f) data related to the performance and effectiveness of their ad purchases.

250. **Opinion No. 13:** To maximize their revenues from the programmatic sale of their display advertising space, and to optimize their auction-related strategies and platform choices, particularly price floors, publishers typically need; (a) data and information regarding the mechanics and rules of the auction process; (b) an understanding of the algorithms employed by the intermediary platforms, tools, and exchanges; (c) data and information about the website visitors who will ultimately receive the display ad; (d) information about the commission, share, take rate, mark-up, or other portion of their payments that is retained by the intermediaries; and (e) data related to the performance and effectiveness of their ad sales.

251. **Opinion No. 14:** Restricting or limiting the availability and flow of necessary and critical information to participants and/or intermediaries in the display ad auction process can affect the transparency and overall fairness of the auction process.

252. **Opinion No. 15:** Similarly, denying symmetrical and fair access to inventory, demand, and functionality to some advertisers, publishers, ad servers, exchanges, or ad buying tools involved in an auction (i.e., unequal distribution of information) can harm the transparency and fairness of the auction process.

253. **Opinion No. 16:** Likewise, changing auction rules without adequate or timely disclosure can harm the transparency and fairness of the auction process.

B. Introduction

254. In light of the centrality of consumer data and information in the digital advertising marketplace, and the profits that can be earned with the help of good quality data and analysis, marketing teams increasingly expend significant resources on data analysis. In this section, I offer some comments on the kinds of informational challenges marketers want to solve with data analysis, the kinds of problems they seek to solve, and the potential rewards for solving those problems.

255. I remark also that I have worked in marketing analysis professionally for twenty-five years—my dissertation is called “Modeling Conversions in Online Advertising”—so I have extensive, direct experience with many of these challenges, and have watched the industry evolve to take on new challenges.

HIGHLY CONFIDENTIAL

256. Finally, publishers and advertisers alike expect transparency when they transact on a programmatic auction.¹²⁹ And Google itself claims to be pro-transparency in display advertising.¹³⁰

257. One of the most basic concepts marketing analysts measure is called efficiency. Efficiency is typically measured as the return on an investment (ROI) or return on ad spend (ROAS). These terms are generally used interchangeably in the advertising world. Marketers are looking for a return on the money they are spending on advertising. When they measure efficiency, they usually divide consumers into three groups: those who have never made a purchase, those who have made one purchase, and those who have purchased more than once.

258. Cost per acquisition is a more specific way of looking at efficiency. Cost per acquisition measures the total cost of a consumer completing a purchase, or how much it costs to get a consumer down the sales funnel from first touch point to

¹²⁹ Deposition of [REDACTED]
144:3.

142:16-

See also, Deposition of
[REDACTED] 153:1–154:9. May 1, 2024.

[REDACTED]
Deposition of
. 62:7–62:11.

¹³⁰ Deposition of [REDACTED] 177:21–178:5, May 21, 2024 177:21 (Q. Is Google, in general, in favor of greater transparency in the ad tech industry? A. Yes, in general, we are. Q. And [REDACTED] you were also, in general, in favor of greater transparency in the ad tech industry, right? A. In general, yes.); Deposition of [REDACTED] 144:1–11, April 30, 2024, (144: 1 Q. Do you think Google should be building towards greater transparency in media buying? A. I think we have a role to play. Q. Is this a comment on Google's transparency? A. This document is regarding a launch that provides greater transparency to media buying in the industry. And so it's something that Google is -- had launched to help provide greater transparency.).

HIGHLY CONFIDENTIAL

conversion.¹³¹ Often the cost per acquisition is offset by a first purchase, and this is something marketers want to know. Knowing the cost per acquisition helps marketers budget, project profit, and optimize future campaigns.

259. Another metric that helps marketers develop marketing strategies is lifetime value. Lifetime value is an estimate of the average revenue a customer will generate during their time as a customer.¹³² When marketers try to improve lifetime value, they need information on repeat customers, and on the cost and ROI for retargeting campaigns.

260. When marketers aim to optimize campaigns and optimize retention specifically, there are a few crucial strategies they adopt, which rely on accurate information. First, there is media optimization, which is when marketers compare the success of campaigns across different media and try to improve the placement of their ads to optimize conversion rates.

261. A second kind of optimization is audience optimization, which is where targeting comes into the picture. When marketers are assessing which kinds of target audiences are most receptive to the products that they are advertising, they look at the effectiveness of different kinds of targeting strategies. The chart below describes several of the targeting strategies that Google makes available.

262. We can consider four types here. The first kind, we can call demographic segments, which allow marketers to advertise to people in a particular demographic category or group of categories, like parents over fifty in Chicago, for example. They can also target based on life events, like the birth of a child. A second type of target is based on the behavior of people in relation to other ads you have placed, like targeting those who have already clicked on an ad, for example. A third way of targeting is called affinity groups, which involves targeting people based on their affinity for a particular interest, like soccer players. A fourth way of targeting is known as exact match, which involves matching email addresses with ads. In this case, some publishers will allow advertisers to upload a set of email lists, and then, if a customer enters that email address the publisher will show an ad to that customer on the site.

263. In the chart below we can see a sample chart of the ways that Google allows advertisers to target audiences.¹³³

¹³¹ Mountain. “What is Cost Per Acquisition (CPA) & How to Calculate it” (undated). Accessed on May 3, 2024. <https://mountain.com/blog/cost-per-acquisition-calculation-how-much-are-you-paying-for-each-customer/>.

¹³² Optimizely. “Lifetime value” (undated). Accessed on June 7, 2024. <https://www.optimizely.com/optimization-glossary/lifetime-value/#:~:text=Lifetime%20Value%20or%20LTV%20is,%2C%20re-sources%2C%20profitability%20and%20forecasting>

¹³³ Word Stream. “Google Ads Audience Targeting: 15 Powerful & Underused Strategies” (January 23, 2024). Accessed on June 4, 2024. <https://www.wordstream.com/blog/ws/2022/09/21/google-ads-audience-targeting-cheat-sheet>.

HIGHLY CONFIDENTIAL

264. In addition to audience optimization and targeting, a third kind of optimization is called creative optimization, which involves measuring the effectiveness of what is called the ad creative, shorthand for the creative work of the advertisement itself. When marketers are considering creative optimization, they are asking about how appealing different ad creatives are for their audiences. Sometimes they will run versions of multiple ads to test appeal for their desired audiences before running a full campaign.

265. Since marketing doesn't happen in a vacuum, and marketers are often competing with other products to land the money and loyalty of customers, they also are in an ongoing competitive position. Marketing strategies evolve and change with the strategies that other competitors adopt. The competitive landscape of marketing is all the more reason that marketers are continually fine tuning and optimizing their data and strategies. The competitive landscape also means that the actions of other marketers and other entities in the digital marketing ecosystem can have a significant bearing on how marketers choose to optimize their campaigns.

266. I briefly describe a handful of marketing measurement techniques that are used to measure the effectiveness of marketing campaigns. In the early days of advertising, marketing effectiveness was primarily measured using Gross Rating Points (GRPs) and Targeted Rating Points (TRPs). GRPs quantify impressions as a percentage of a total population, typically the entire US, and provide an estimate of the audience reach. TRPs are more refined, focusing on specific demographic targets. These metrics were essential in traditional media like television and radio, offering a broad gauge of campaign reach and frequency.

267. The most rigorous traditional media channel is direct mail. Direct mail has a long history of using experimentation to measure effectiveness. Marketers

HIGHLY CONFIDENTIAL

would send different versions of mail to various segments and analyze the responses to determine the most effective strategies. This method laid the groundwork for modern experimental techniques by emphasizing the importance of data-driven decisions and controlled testing environments.

268. As marketing evolved, Marketing Mix Modeling (MMM) emerged as a scientific approach to evaluating the impact of marketing channels on sales and conversions. MMM uses statistical analysis to assess the performance of different media types, allowing marketers to allocate budgets more effectively across channels. This approach helps in understanding how different marketing efforts contribute to overall business goals, providing a more holistic view of campaign effectiveness. MMM is not a perfect approach, as there are aspects of consumer behavior that cannot be included in the model.¹³⁴ Nevertheless, MMMs give marketers the ability to shift spend between channels, optimize their media mix, and forecast future performance.

269. While MMMs are rigorous, it is important to point out that practitioners consider the utilization of these models as much art as science. A famous quote in marketing measurement, attributed to the department store owner John Wanamaker, is the practitioner's touchstone: "I know half the money I spend on advertising is wasted; the trouble is I don't know which half."¹³⁵ There is a great deal of statistical noise in the marketing measurement system, and discerning the cause of changes to performance can be difficult. This causes marketers to adopt a "don't rock the boat" attitude toward shifting the marketing mix. The company Measured, a leading experimentation platform, recommends taking the insights from MMM with a grain of salt unless supported by randomized control trials.¹³⁶ Making small shifts can make measurement more difficult, but the practice is common enough that Zhou, et al state, "Variation in marketing allocations over time can be insufficient to produce well-identified estimates of the effects of different marketing actions."¹³⁷

270. Marketing budgets at large advertisers are shifted on several different time scales. Typically, companies start by setting an overall annual marketing budget. This budget is based on past performance, projected revenues, and strategic

¹³⁴ MMM suffers from an "exogeneity problem". For instance, models alone cannot estimate the ROI of branded search, where a consumer searches for very specific terms such as "toyota dealers sacramento", since those consumers are close to the bottom of the funnel.

¹³⁵ B2B Marketing. "Half the Money I Spend on Advertising is Wasted; The Trouble is I Don't Know Which Half" (undated). Accessed on June 4, 2024. <https://www.b2bmarketing.net/archive/half-the-money-i-spend-on-advertising-is-wasted-the-trouble-is-i-dont-know-which-half-b2b-marketing-2/>.

¹³⁶ Measured. "The Beginner's Guide to Marketing Mix Modeling (MMM)" (February 8, 2023). Accessed on June 4, 2024. <https://www.measured.com/blog/the-beginners-guide-to-marketing-mix-modeling-mmm/>.

¹³⁷ Zhou, G., Skokan, I., Runge, J. "Packaging Up Media Mix Modeling: An Introduction to Robyn's Open-Source Approach." 2024.

HIGHLY CONFIDENTIAL

priorities for the upcoming year.¹³⁸ The total dollar amount is allocated across various marketing channels and activities. Once these allocations are made, spends rarely shift across channels dramatically and, as McKinsey recommends,

Don't go cold turkey. A drastic budget shift in one year could complicate vendor relationships and marketing activities. Your organization may have product-launches or campaigns it cannot alter. Budgeting shifts can and should be phased in through pilot programs that offer early evidence of success and learnings along the way.¹³⁹

Reallocating budget across channels is more similar to turning an oil tanker than driving a race car. Within channels, where budgets are managed within a smaller team, changes can be more dramatic and happen as often as week to week.

271. In the past, marketers conducted experiments by targeting specific regions of the country to test different strategies. Today, advanced techniques like Randomized Controlled Trials (RCTs) and Ghost Ads¹⁴⁰ have revolutionized experimentation. RCTs involve randomly assigning subjects to different groups to isolate the effects of marketing interventions. Ghost Ads, on the other hand, track the impact of ads that are planned but not shown to a control group, providing insights into the true effectiveness of digital campaigns. These modern methods allow for more precise and scalable measurements, enabling marketers to optimize their strategies with greater confidence.

C. Tracking Activity Online

272. Cookies are a set of randomly generated numbers on your computer that identify your browser. Cookies are unique to different browsers and different computers. So, for example, I have different cookies for Chrome and Firefox. I also have different cookies on different computers I work on. Because cookies are a unique identifier, they allow someone to recognize that a page or part of a page is served to me. There is an opportunity, in other words, for the browser to read the cookie and report it back to an ad server—whether a publisher or an advertiser—and either of them knows that the browser hit that page. When data analysts look at the views on a page, they are engaged in what is called view measurement. View measurement helps marketing analysts determine whether a person saw an ad. When an ad is viewed, the ad server creates a log that lists viewers' engagement with an ad.

¹³⁸ Hubspot. "How to Manage Your Entire Marketing Budget [Free Budget Planner Templates]" (August 12, 2022). Accessed on June 4, 2024. <https://blog.hubspot.com/marketing/how-to-manage-marketing-budget-free-budget-templates>.

¹³⁹ McKinsey & Company. "How to Reallocate Marketing Budgets to Drive Growth" (April 1, 2015). Accessed on June 4, 2024. <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/how-to-reallocate-marketing-budgets-to-drive-growth>.

¹⁴⁰ Johnson, G., Lewis, R., and Nubbemeyer, E. "Ghost Ads: Improving the Economics of Measuring Online Ad Effectiveness" *Sage Journals* vol. 54, no. 6. 2017.

HIGHLY CONFIDENTIAL

273. When a user sees an ad on a website, the ad server sets a cookie on the user's browser. Cookies can be categorized into first-party cookies, which are set by the website the user is directly visiting, and third-party cookies, which are set by domains other than the one the user is visiting. Third-party cookies are commonly used by ad tech platforms to track user behavior across different websites and build a comprehensive profile of the user's interests and activities.

274. In a single display advertising campaign, multiple entities might be involved in serving and tracking the ads, leading a consumer to end up with several cookies on their device from different sources. For example, when a user visits a website that displays ads, the site places a cookie on the user's browser to track which ads have been shown to the user and how the user interacts with those ads. If the ad is served through a demand-side platform (DSP) like DV360, the DSP also places its own cookie on the user's browser to manage bids, track ad impressions, and measure campaign performance.

275. Additionally, the campaign might involve third-party data providers that supply targeting data. These providers place their own cookies to track user behavior and supplement the targeting information with their data. If the campaign includes retargeting strategies, retargeting platforms will place their own cookies to identify users who have previously visited the advertiser's website or engaged with their ads, enabling the platform to show relevant ads to users as they browse other websites. Furthermore, the website (publisher) where the ad is displayed may also set cookies to track user interactions with their content and ads, optimizing ad placements and enhancing user experience.

276. In short, a single advertising impression can generate or require the reading of many cookies on the consumer's computer.

277. Cookie matching is a process used in digital advertising to synchronize the unique identifiers assigned to users by different ad tech platforms. When users visit various websites, each ad tech platform involved in displaying ads sets its own cookie on the user's browser, each containing a unique identifier. Since these identifiers are platform-specific, cookie matching enables these platforms to recognize the same user across different websites and devices, facilitating more effective tracking, targeting, and measurement of ad campaigns.

278. The cookie matching process typically involves multiple redirects and data exchanges between ad tech platforms. For example, when a user visits a website with an ad tag from Platform A, this platform reads its cookie and sends a request to Platform B's server, including the user's unique identifier from Platform A. Platform B then sets its own cookie and links it to Platform A's identifier, establishing a connection between the two. This synchronization allows both platforms to share information about the user's behavior and preferences, enabling them to deliver more relevant and personalized ads.

HIGHLY CONFIDENTIAL

279. While cookie matching is a powerful tool for enhancing the effectiveness of programmatic advertising, it comes with several challenges. One of the primary difficulties is the inherent complexity of the process itself. Synchronizing cookies between different platforms involves multiple redirects and data exchanges, which can introduce latency and reduce the accuracy of matches. Each redirect increases the time it takes to load the webpage, potentially degrading the user experience and leading to higher bounce rates. I have been involved with projects to do this cookie matching in a parallel process, so that the matching minimizes the delay in serving the web page to the user.

280. Another significant challenge is the impact of user behavior and browser settings on cookie persistence. Users often clear their cookies or use browser settings that block third-party cookies, which can disrupt the matching process. When cookies are cleared, the unique identifiers assigned by the ad tech platforms are lost, requiring the platforms to start the tracking and matching process anew. Similarly, privacy-focused browser settings and extensions, such as those that block third-party cookies, can prevent ad tech platforms from setting and reading cookies, thereby hindering their ability to synchronize user identifiers. Parties such as Google, who maintain persistent ideas via information like email addresses, have a significant edge in surmounting this problem.

281. The growing emphasis on privacy and data protection also poses challenges to cookie matching. Regulations such as the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) in the United States impose strict requirements on how user data can be collected, stored, and used. Compliance with these regulations often necessitates obtaining explicit user consent before setting cookies, further complicating the cookie matching process.

282. Finally, the technical limitations of cookie-based tracking, such as the inability to track users across different devices reliably, present additional hurdles. As users increasingly engage with digital content across multiple devices, achieving a seamless cross-device tracking and matching becomes crucial yet challenging. Once again, companies such as Google who maintain an extensive desktop and smartphone presence have advantages in handling the cross-device problem.

283. Despite these challenges, cookie matching remains a cornerstone of digital advertising, enabling more precise targeting, measurement, and optimization of ad campaigns.

284. Google's scale advantage in the digital advertising ecosystem is formidable, primarily due to its vast reach and extensive first-party data. Google's numerous services, including Google Search, YouTube, Google Maps, Gmail, and the Android operating system, allow it to collect a wealth of user data across multiple touchpoints. This data provides Google with unparalleled insights into user behavior, preferences, and interests, enabling highly accurate targeting and personalization of ads.

HIGHLY CONFIDENTIAL

Additionally, Google's widespread presence on publisher sites through tools like Google Analytics and Google Ad Manager ensures that its tags are embedded across a vast array of web properties, facilitating superior cookie matching and higher match rates.

285. Competing against Google's scale requires other players in the ad tech space to leverage several key components. They must establish robust data partnerships to enrich their third-party data and enhance targeting capabilities. By collaborating with data providers and integrating data from various sources, these platforms can build comprehensive user profiles that rival those created with Google's first-party data. Additionally, they must develop advanced algorithms and machine learning techniques for probabilistic matching, which can improve match rates and targeting accuracy even without extensive first-party data. As I discuss below and in Section X, Google's "Privacy Sandbox" program will undercut these efforts.

D. Google's Privacy Sandbox Program

286. As discussed above, cookies are unique identifiers maintained by web browsers that are assigned to each browser and machine.¹⁴¹ Cookies facilitate targeted display advertising by providing a mechanism for display advertising tools to identify and track individual users and their browsing history so that relevant advertisements may be placed during the user's future web visits based on their past browsing history. Both Google and its competitors rely on this data to place display advertisements with specific users and without the ability to access cookie data, advertising tools can only guess where to display advertisements.¹⁴²

287. Google has a unique position relative to third-party cookie data in particular because its Chrome Internet browser is used by over 3 billion users (over 60% of users).¹⁴³

[REDACTED]¹⁴⁴ Because Chrome is so widely used and one of the largest available source of third-party cookie data, much of the data needed for targeting by display advertising tools flows through Chrome. Many third-party publishers and advertisers [REDACTED]

[REDACTED]—have expressed

¹⁴¹ See <https://developers.google.com/tag-platform/security/concepts/cookies#:~:text=Cookies%20are%20small%20files%20saved,ads%20or%20measure%20your%20success>.

¹⁴² See, e.g., <https://business.safety.google/adscookies/> ("Google uses cookies for advertising, including serving and rendering ads, personalizing ads (depending on your ad settings at g.co/adsettings), limiting the number of times an ad is shown to a user, muting ads you have chosen to stop seeing, and measuring the effectiveness of ads.").

¹⁴³ Daniel Taylor Dep. 253:21-254:16; <https://backlinko.com/chrome-users>; Fleck, Anna. "Google's Chrome Has Taken Over the World." Digital image. September 1, 2023. Accessed June 07, 2024. <https://www.statista.com/chart/30734/browser-market-share-by-region/>.

¹⁴⁴ Fleck, Anna. "Google's Chrome Has Taken Over the World." Digital image. September 1, 2023. Accessed June 07, 2024. [https://www.statista.com/chart/30734/browser-market-share-by-region/](https://www.statista.com/chart/30734/browser-market-share-by-region/;); [REDACTED]

HIGHLY CONFIDENTIAL

that their access to cookie data from Chrome is important to their ability to advertise effectively.¹⁴⁵ While third-party cookie data does raise certain privacy concerns, the benefits include:¹⁴⁶

- Personalized advertising: They can enable advertisers to personalize ads based on a user's browsing history, interests, and demographics, making ads more relevant to the user.
- Better user experience: They[] can help websites remember user preferences and login information, which can improve the user experience.
- User analytics: They can help website owners understand how users are interacting with their site, which can help them to improve its performance.

288. Google is currently discussing the deprecation of third-party cookie data in Google Chrome and replacing it with a Google program called "Privacy Sandbox," which is described and discussed in Plaintiffs' Fourth Amended Complaint in this action at 164-168, ¶¶ 473-481.

289. [REDACTED] described Privacy Sandbox as including "an intention to remove or block third-party cookies from being available in its browser."¹⁴⁷ [REDACTED] explained that "third-party cookies are a piece of technology that is used by many websites as well as ad tech companies to understand user behavior across websites" and, therefore, "blocking those cookies would then impact those technologies that use the cookies."¹⁴⁸ In place of third-party cookies used to track and measure users across websites, "Chrome, who is overseeing the Privacy Sandbox initiative, is proposing a set of APIs, which stands for application programming interface, to provide to ad technology providers to help them deliver ads and measure the performance of them without the need for third-party cookies and the associated tracking that comes with it."¹⁴⁹ Let me put a fine point on this API comment: Google is proposing that third-party cookies, an egalitarian approach to user measurement, be replaced by a series of server calls with Google as the intermediary.

290. In addition, according to [REDACTED] "historically, technology companies would use third-party cookies in Chrome -- set in Chrome's browsers, or set on websites that are accessible via Chrome's browsers, to track people's activity across the web to understand what they're interested in, to measure outcomes, things like that," and that the "Sandbox APIs are Chrome's proposed solution to give access to information about users on Chrome that can help them deliver those use cases without sharing individual user information."¹⁵⁰ Google will offer access to the data that it has

¹⁴⁵ See [REDACTED]

[REDACTED]).

¹⁴⁶ <https://www.cookieyes.com/blog/third-party-cookies/>.

¹⁴⁷ Deposition of [REDACTED] 249:5-9.

¹⁴⁸ *Id.* at 249:13-20.

¹⁴⁹ *Id.* at 250:2-9.

¹⁵⁰ *Id.* at 252:6-17.

HIGHLY CONFIDENTIAL

collected on its Chrome users through those APIs.¹⁵¹ Google will also take steps to anonymize and group user data.¹⁵² Google will also gain a measure of control in the ad tech stack they have heretofore not enjoyed.

291. Google's deprecation of third-party cookie data is concerning to competing display advertising tools and to publishers and advertisers for several reasons. First, by deprecating third-party access to cookie data, Google is acting as a gatekeeper to information that is necessary to effectively compete with Google's advertising tools and thereby protecting the market position of Google's various advertising tools. In other words, depreciation of cookies exacerbates the problem of unequal access to information between Google and its competitors. This is especially true in view of Google's unique access to cookie data from its own Chrome product. Google may use that data to benefit its own advertising tools while excluding access to any potential competitors that might use that same data to compete with Google's advertising tools. The UK Competition and Markets Authority (CMA) reported that this was a concern in its recent Q1 2024 report, stating that the CMA seeks to ensure "that Google does not design, develop or use the Privacy Sandbox tools in ways that reinforce the existing market position of its advertising products and services, including Google Ad Manager (GAM)."¹⁵³

292. Second, efforts to deprecate third-party access to cookies, such as Google's Privacy Sandbox, create a conflict of interest. On the one hand, Google represents the interests of publisher tools, advertiser tools, and exchanges in creating its Privacy Sandbox, Audiences, Interest Groups, and other similar features to replace access to cookie data. But at the same time, Google has competing products that incentivize Google to either limit the capabilities of the features that replace access to cookie data and/or for Google to grant itself advantaged access to Chrome cookie data or the Sandbox features it has created so that Google's products outperform competing advertising tools. The UK CMA shares this concern and has flagged that Sandbox includes a risk that Google will "self-preference...through their design development or implementation."¹⁵⁴

¹⁵¹ *Id.* at 251:5-24.

¹⁵² See <https://iabtechlab.com/wp-content/uploads/2024/02/Privacy-Sandbox-Fit-Gap-Analysis-PUBLIC-COMMENT-RELEASE.pdf> at 20-22 (discussing Sandbox's use of "Audiences" and "Interest Groups" to enable targeted advertising without sharing user data with advertising tools).

¹⁵³ CMA Q1 2024 update report on implementation of the Privacy Sandbox commitment at ¶14(a) (available at (https://assets.publishing.service.gov.uk/media/662baa3efee48e2ee6b81eb1/1._CMA_Q1_2024_update_report_on_Google_Privacy_Sandbox_commitments.pdf); see also ¶¶140-141 ("We are aware that there are questions and concerns around Google's use of data" and enumerating a number of open questions regarding how Google will access browser data while limiting third party access via its Sandbox.).

¹⁵⁴ *Id.* at ¶131(c); see also ¶158(c) (accusing Google of having already "self-preferencing its advertising products and services, and using competitively sensitive information provided by an ad tech provider or publisher to Chrome for a purpose other than that for which it was provided" and also demonstrating a tendency for Google to self-preference in contradiction of assurances that it will not do so.).

HIGHLY CONFIDENTIAL

293. Third, Google's Privacy Sandbox solutions currently contemplate moving correlation of user information to the browser that has traditionally been done in AdTech platforms. For example, Google proposed a feature it calls TURTLEDOVE where “[t]he AdTech platform won't know that [two ad requests are coming from...to make it hard for AdTech platforms to identify users by connecting the time the two requests are sent.”¹⁵⁵ Consequently, “many of the key ad-decisioning and even auction mechanics will be conducted in the browser (aka on device) instead of by AdTech platform.”¹⁵⁶ By moving functionality from AdTech products to its own browser, Google limits competing products' ability to analyze raw user data.

E. Measuring Marketing Effectiveness

294. Many marketing decisions are made using much less sophisticated approaches. Key metrics for the basics of marketing measurement include Click-Through Rate (CTR), conversion rate, and attributed last ad conversions.

295. If someone clicks on an ad, there is a click through URL that takes that user to another page. While the user is clicking on a page, the browser is tracking the movement and logs each step. In the early days of display marketing, views were measured in terms of the number of ads served. In the mid 1990s, view performance was measured in click throughs. For several years, marketers worked to optimize media placement by increasing clicks.

296. Click-Through Rate (CTR) is a fundamental metric that measures the ratio of users who click on an ad to the number of total users who view the ad (impressions). A higher CTR indicates that the ad is effective in attracting interest and prompting action. It is calculated as:

$$CTR = \frac{Number\ of\ Clicks}{Number\ of\ Impressions} \cdot 100$$

CTR helps marketers gauge the relevance and appeal of their ad content and creative elements. It provides, however, only a partial picture since the people who click on an ad may immediately navigate away or simply fail to purchase.

297. Conversion Rate goes a step further by measuring the percentage of users who complete a desired action (such as making a purchase, signing up for a newsletter, or filling out a form) after clicking on an ad. This metric is essential for assessing the effectiveness of the ad in driving valuable actions. It is calculated as:

$$Conversion\ Rate = \frac{Number\ of\ Conversions}{Number\ of\ Impressions} \cdot 100$$

¹⁵⁵ Mike Sweeney, How Google Chrome's Privacy Sandbox Will Work + Possible Solutions for AdTech (May 15, 2024) (available at <https://clearcode.cc/blog/chrome-privacy-sandbox-explained/#toc-label-7>).

¹⁵⁶ *Id.*

HIGHLY CONFIDENTIAL

A high conversion rate indicates that the ad is not only attracting clicks but also leading to meaningful interactions and transactions.

298. The previous equation for conversion rate includes the metric *conversions*, the number of desired actions undertaken. But this metric does not exist in isolation, and we need additional information to calculate it. In the most simplistic form, we can simply count the number of actions in a certain time period. This approach fails to tie the advertising to the action, however. The missing step is called “attribution”, where the advertising is tied to the action.¹⁵⁷

299. The most simplistic approach to attribution is called last ad attribution. When I began my career, I was part of the first analytics team to make conversions available to advertisers. Last ad conversions refer to conversions where 100% of the credit for a conversion is given to the final ad interactions that directly precede a conversion. This metric is fraught—many consumers may see multiple advertisements and some channels, e.g., branded search, are much more likely to be the last advertisement a converter sees. Nevertheless, by analyzing last ad conversions, marketers can identify which ads and channels contribute most significantly to the final conversion, enabling more informed allocation of advertising budgets.

300. Eventually, they also began to look at view-through conversion tracking, which tracked users when they saw an ad and didn’t click on it, but later came back and made a purchase.

301. In addition to these primary metrics, marketers often track other key performance indicators (KPIs) such as ROI, Cost Per Acquisition (CPA), and Customer Lifetime Value (CLV). ROI measures the revenue generated for every dollar spent on advertising, helping to assess the financial efficiency of ad campaigns. CPA tracks the cost associated with acquiring a new customer, while CLV estimates the total revenue a business can expect from a customer over the course of their relationship.

302. Together, these metrics provide a comprehensive view of marketing performance, allowing businesses to make data-driven decisions to optimize their campaigns and maximize ROI. By continuously monitoring and analyzing these metrics, marketers can refine their strategies to better meet their objectives and drive sustained growth.

303. I will speak more about auctions and Real-time Bidding (RTB) in the following section. At this point, let me say that RTB allows advertisers to target specific users based on detailed bidstream data, which includes sensitive information

¹⁵⁷ Generally, in advertising the term “action” refers to a particular behavior, such as purchasing, visiting a certain section of a site, or signing up for an email newsletter. The term “conversion” refers to actions that have been associated with advertising, reflecting the fact that someone who was reached by an advertisement has now been converted to a customer.

HIGHLY CONFIDENTIAL

such as IP addresses and browsing history.¹⁵⁸ Google has collected data and stored it, often without the explicit consent of the user, raising significant privacy concerns around the world.¹⁵⁹ Despite these issues, RTB has grown rapidly, driven by its ability to deliver highly targeted advertising at scale.¹⁶⁰

304. Multi-touch attribution model (MTA) is a sophisticated approach to marketing analytics that aims to provide a comprehensive understanding of the effectiveness of various marketing touchpoints in driving conversions.¹⁶¹ Unlike traditional single-point attribution models, like the last ad model, which attribute full credit to a single touchpoint encountered before a conversion, MTA was introduced as a way to share credit across touchpoints. Using MTA enables advertisers to understand each touchpoint's role in creating a customer.¹⁶²

305. At its core, MTA acknowledges the complexity of contemporary consumer behavior and recognizes that multiple interactions with a brand across various channels can influence a customer's decision to convert. By evaluating and weighing the impact of several touchpoints, MTA provides marketers with valuable insights into which marketing activities contribute most effectively to the bottom line.¹⁶³ This enables them to allocate their budgets more efficiently by investing in touchpoints that yield the highest returns while minimizing spending on less effective ones.

F. Summarizing the Data in the Auction Process

306. The previous subsection discussed marketing measurement generally. In this subsection I will outline the data available to typical participants in a programmatic display auction.

307. Programmatic display auctions are complex, data-driven processes that provide various types of information to participants, including advertisers, publishers, and intermediaries like DSPs, SSPs, and exchanges. Understanding the data available at each stage of the auction process is essential for understanding the

¹⁵⁸ Forbes. “Real-Time Bidding: The Ad Industry Has Crossed a Very Dangerous Line” (October 18, 2021). Accessed on June 4, 2024. <https://www.forbes.com/sites/hessiejones/2021/10/18/real-time-bidding-the-ad-industry-has-crossed-a-very-dangerous-line/?sh=6dd0258148ca>.

¹⁵⁹ Reuters. “Google settles \$5 billion consumer privacy lawsuit” (December 29, 2023). Accessed on June 5, 2024. <https://www.reuters.com/legal/google-settles-5-billion-consumer-privacy-lawsuit-2023-12-28/>; See also BBC. “Google settles \$5bn lawsuit for ‘private mode’ tracking” (December 28, 2023) Accessed on June 5, 2024. <https://www.bbc.co.uk/news/business-67838384>.

¹⁶⁰ Forbes. “Real-Time Bidding: The Ad Industry Has Crossed a Very Dangerous Line” (October 18, 2021). Accessed on June 4, 2024. <https://www.forbes.com/sites/hessiejones/2021/10/18/real-time-bidding-the-ad-industry-has-crossed-a-very-dangerous-line/?sh=6dd0258148ca>.

¹⁶¹ Search Engine Journal. “Complete Guide to B2B Multitouch Attribution Models” (June 2, 2022). Accessed on June 4, 2024. <https://www.searchenginejournal.com/b2b-multitouch-attribution-models/450898/>.

¹⁶² *Id.*

¹⁶³ *Id.*

HIGHLY CONFIDENTIAL

unique position that Google occupies in this space and then the ways in which they took advantage of that position.

308. Advertisers are the buy side in the programmatic auction process, leveraging their unique data to inform their bidding strategies. They typically use user data, which includes demographic information, interests, and online behavior, to target their desired audience effectively. On the buy side, the scale of the data a firm possesses, the quality of that data, and the recency of that data are all competitive advantages. Advertisers also consider ad placement data, encompassing details about the context, size, and format of the ad. Performance metrics from historical campaigns, such as CTRs and conversion rates, provide insights into the potential effectiveness of their bids. Additionally, understanding auction dynamics, including estimates of bid floor prices and competition levels, helps advertisers optimize their bids to achieve maximum impact and ROI.

309. Publishers represent the supply side of the programmatic auction process and provide the ad inventory and user audience that advertisers seek to engage. They rely on detailed audience insights, including demographic and behavioral data, to attract high-quality bids and to package their inventory in an attractive fashion. Ad performance data, such as viewability and engagement rates, helps publishers understand which types of ads resonate best with their audience. Access to this performance data can help a publisher decide whether they want to send inventory into indirect sales channels like exchanges or networks. Alternatively, they can decide to retain control over that inventory and sell it in the direct channels, outlined in Section VI.C. Inventory availability data, including real-time forecasts and fill rates, allows publishers to manage their ad space efficiently. By analyzing bid data, publishers can adjust their pricing strategies to maximize revenue, ensuring they attract competitive bids while maintaining a positive user experience. Publishers are also responsible for setting the price or bid floors for their inventory.

310. Price floors are critically important for publishers. While I was at Microsoft I spent months building models to create these floors in the nascent programmatic space. Imagine you work for a publisher like MSN or Yahoo!, where your inventory can be organized into content verticals such as finance, sports, news, and astrology. On average, these verticals will garner different CPMs in direct deals and the goal with programmatic is to sell your remnant¹⁶⁴ for as much money as possible, attempting to approach the CPMs of the direct deals. The rate that you are able to sell your inventory at is called your “effective CPM” or eCPM. Typically, some verticals, such as finance, command relatively high prices, while others, such as astrology or games, sell for lower prices.

311. But programmatic inventory is sold on an impression-by-impression basis, not in large blocks like direct deals. And so, the eCPM of your astrology vertical

¹⁶⁴ The industry term for unsold inventory is “remnant”.

HIGHLY CONFIDENTIAL

is an average across all the impressions to all of the users you see on that site. Since some wealthy people who are interested in finance are also interested in astrology, an information battle begins. An advertiser who knows that a given astrology impression is being delivered to a finance aficionado could theoretically pick up that impression for a much lower rate than doing a direct deal on the finance vertical. This is where price floors come in. Publishers can identify the high-value impressions regardless of where they appear on the site and set price floors that protect their best inventory, ensuring adequate revenue and protecting the prices of their direct deals.

312. Advertiser ad buying tools act as intermediaries that help advertisers manage and optimize their bidding strategies in real-time. These tools aggregate market data, providing advertisers with insights into pricing trends, demand fluctuations, and competitive benchmarks. They can facilitate access to multiple ad exchanges, enabling advertisers to participate in numerous auctions simultaneously. These tools provide real-time auction metrics, such as bid response times and win rates, to help advertisers fine-tune their bids. Additionally, advertiser ad buying tools ensure technical performance by minimizing latency and error rates, thereby maximizing the efficiency and effectiveness of ad delivery.

313. Similarly, publisher ad servers work on behalf of publishers to optimize the sale of their ad inventory. These products provide publishers with market data and analytics to help them understand demand trends and price their inventory competitively. They facilitate real-time auctions, connecting publishers with demand sources. By offering advanced yield optimization tools, publisher ad servers help publishers maximize their revenue from ad sales. These servers also ensure the technical performance of ad delivery, managing factors like latency and error rates to maintain a smooth and efficient auction process.

314. Ad exchanges are digital marketplaces where publishers and advertisers come together to buy and sell ad inventory in real time. These exchanges facilitate programmatic auctions, allowing advertisers to bid on individual ad impressions based on the data provided by the buy side and the sell side. Exchanges aggregate supply and demand, providing a platform for efficient, competitive, and, theoretically, fair bidding. Exchanges should play a critical role in ensuring transparency and fairness in the auction process, offering insights into bid prices, win rates, and market trends. By enabling real-time transactions, ad exchanges should help both publishers and advertisers achieve their objectives in the digital advertising ecosystem.

315. The exchange take rate, or commission, is the percentage of the transaction value that an ad exchange retains from the auction. This fee is deducted from the winning bid before the remaining amount is paid to the publisher. If the take rate is too high, more bids will drop below the publisher's price floor, causing the auction to fail to clear. This situation leads to unsold inventory, reducing revenue for both publishers and exchanges. Therefore, maintaining a balanced take rate is crucial for ensuring a healthy and competitive auction environment.

HIGHLY CONFIDENTIAL

G. The Potential Effects of Information Imbalances, Exclusions, Failure to Disclose, and Misrepresentations

316. In Section VI.D I discuss the various features of programmatic auctions and, in particular, the way auction participants think about sealed-bid second-price auctions. In this Section IX I addressed the great lengths that industry practitioners go to in the process of acquiring and leveraging this information. Section VIII explains Google's unique position in the ad tech ecosystem. In Section X I look more carefully at how Google exploits this position, both in terms of its market position in the largest sectors of the ad tech market, and in terms of its unique access to information. I conclude that Google has taken steps to increase earnings by taking advantage of its informational access, and that it has stifled competition using its multifaceted market position.

317. I conclude this section in general terms, describing information that typical exchange participants do *not* have and what having that information would mean.

1. The Buy Side Wish List

318. In an exchange, the buy side is seeking to purchase high-quality inventory at the lowest prices possible. This subsection explores the types of information that the buy side might wish it could have. This subsection explores counterfactuals and hypotheticals. Section X holds my analysis of the ways in which Google exploited its unique market position to take advantage of some of these scenarios.

319. Advertisers have a limited view of the inventory and their competition. With additional data they could optimize their campaigns and enhance their bidding strategies. One critical piece of information they would like is detailed user engagement data, which includes metrics on how users interact with ads on a given publisher's site, such as time spent on the ad, click patterns, and overall engagement rates. This data helps advertisers understand the effectiveness of their ads in capturing user interest and can inform adjustments to creative elements and targeting strategies. This data is only available to participants on the sell side.

320. Another valuable insight for advertisers would be viewability metrics, which provide information on the likelihood that an ad will be seen by users. Understanding viewability rates for specific placements allows advertisers to focus their bids on inventory with higher visibility, ensuring that their ads are seen by the intended audience and not hidden in less visible parts of a webpage. This data is typically only available to participants on the sell side.

321. Fraud detection information would also be crucial for advertisers. Access to data on the measures in place to detect and prevent ad fraud would ensure that their ads are served to real, engaged users, protecting their investment and maximizing the return on their ad spend. This includes knowing the prevalence of bot traffic

HIGHLY CONFIDENTIAL

and other fraudulent activities on the platforms they are bidding on. This information is generally only available on the sell side.

322. Advertisers would also benefit from detailed ad placement context, including information about the content and environment in which their ads will appear. This ensures brand safety and relevance, as advertisers can avoid placing their ads next to inappropriate or irrelevant content, which could negatively impact brand perception. This information is generally only available on the sell side.

323. Access to historical performance data from previous similar auctions would another key piece of information for advertisers. While advertisers know the winning price for auctions they win, they do not know the winning prices or the price floors on auctions they lose. This data would include winning bid ranges, performance metrics, and trends from past auctions, allowing advertisers to benchmark their bids and strategies against historical norms. By understanding past auction dynamics, advertisers can better predict and strategize for future auctions, enhancing their chances of winning valuable ad placements at optimal prices. This information in totality is only available to exchange participants, though the sell side has some of this.

324. Advertisers would love to know the price floors on inventory they are bidding on. With that information they could adopt the optimal “Price is Right” strategy of bidding one cent more than the floor for auctions they were going to win, ensuring optimal return on the impression. This information is available only to the sell side and the auction.

325. In summary, advertisers would benefit greatly from having access to user engagement data, viewability metrics, fraud detection information, ad placement context, and historical performance data. These insights would enable them to refine their bidding strategies, ensure their ads are seen by the right audience, and achieve better campaign outcomes in the competitive landscape of programmatic auctions.

2. The Sell Side Wish List

326. In an exchange, the sell side is seeking to maximize its total revenue by selling as many impressions as possible for as high a price as possible. This subsection explores the types of information that the sell side might wish it could have. This subsection explores counterfactuals and hypotheticals. Section X holds my analysis of the ways in which Google exploited its unique market position to take advantage of some of these scenarios.

327. Publishers play a critical role in the programmatic auction process by providing the ad inventory and audience data that attract advertisers. They would greatly benefit from having visibility into all bids submitted during auctions. This information would enable them to set optimal price floors that protect their revenue,

HIGHLY CONFIDENTIAL

ensuring high-value impressions are not undersold. By knowing the range of bids, publishers could adjust their price floors to capture more value from their inventory. This information is only known by the exchange.

328. Knowledge of losing bids has other benefits for publishers. Understanding the strategies and bid amounts from advertisers that did not win can help publishers gauge market demand more accurately. This insight allows them to adjust their inventory management and pricing models, ensuring they remain competitive while maximizing revenue. This information is only known by the exchange.

329. Advertiser data would also be valuable for publishers. Detailed insights into advertiser behavior and preferences could enable publishers to package and market their inventory more attractively. By understanding which advertisers are bidding and their specific interests, publishers could tailor their offerings to meet advertiser needs better, potentially increasing bid amounts and revenue. This information is only known by the buy side.

330. Additionally, publishers would benefit from having access to real-time market trends and performance metrics. This data would include information on current demand trends, pricing benchmarks, and ad performance metrics such as viewability and engagement rates. With these insights, publishers could optimize their inventory allocation, improve ad placements, and adjust their strategies to maximize revenue. This information is partially known by exchanges and by the buy side.

331. In summary, publishers would greatly benefit from access to all bids submitted during auctions, knowledge of losing bids, detailed advertiser data, and real-time market trends. These insights would enable them to refine their pricing strategies, optimize their inventory management, and achieve better revenue outcomes in the competitive landscape of programmatic auctions.

3. The Exchange Wish List

332. Ad exchanges serve as the digital marketplaces where publishers and advertisers come together to buy and sell ad inventory in real-time. This subsection explores the types of information that exchanges could exploit with access to data that exchanges typically do not have or if the exchanges were willing to manipulate their auctions. This subsection explores counterfactuals and hypotheticals. Section X holds my analysis of the ways in which Google exploited its unique market position to take advantage of some of these scenarios.

333. There are several ways exchanges could manipulate the auction process to their advantage.

334. One way that exchanges could manipulate auctions is to implement a “right of first refusal” for their own bids or bids from preferred partners. This would allow the exchange to optimize bidding strategy without having to worry about

HIGHLY CONFIDENTIAL

competitive bids, while disadvantaging advertisers and publishers. Advertisers would be harmed because valuable inventory could be “cherry picked” before they had an opportunity to bid on it. Publishers would be harmed because some higher bids than the ultimate clearing price would not be “opened”.

335. Another way that exchanges could manipulate auctions is by having a “last look” at incoming bids. This practice would allow exchanges to see the highest bid and adjust their own bids or bids from preferred partners to just outbid the highest offer, ensuring their inventory wins the auction. This could undermine the fairness of the auction and skew results in favor of certain participants.

336. Exchanges could also adjust their take rates dynamically based on the incoming bids. By increasing the take rate when bids are high and decreasing it when bids are low, exchanges could maximize their revenue while ensuring auctions clear. This manipulation could lead to inefficiencies and reduced transparency, as advertisers and publishers would not have a clear understanding of the actual costs involved.

337. Another unethical practice would be prioritizing certain bids over others based on undisclosed criteria. Exchanges could give preferential treatment to bids from partners with whom they have special agreements, regardless of the bid amounts. This could distort the market, making it difficult for other advertisers to compete fairly.

338. Exchanges could also manipulate the auction mechanics, such as altering the bid increments or changing the auction type from a second-price to a first-price auction without informing participants. This could lead to higher costs for advertisers who are unaware of the changes, ultimately benefiting the exchange financially.

339. Additionally, exchanges could withhold critical market data from participants. By not sharing information about bid landscapes, losing bids, and auction outcomes, exchanges could keep advertisers and publishers in the dark, making it challenging for them to optimize their strategies effectively.

340. In sum, if exchanges were willing to be unethical, they could manipulate the auction process through practices like last-look bidding, dynamic take rate adjustments, bid prioritization, altering auction mechanics, and withholding market data. These manipulations would undermine the transparency and fairness of the auction process, disadvantaging both advertisers and publishers.

X. An Analysis of Google’s Activities in the Website Display Ecosystem

A. Opinion 17

341. **Opinion No. 17:** Google’s Bernanke, Global Bernanke, Bell, Reserve Price Optimization (RPO), Dynamic Revenue Sharing (DRS), Poirot, Elmo, Exchange

HIGHLY CONFIDENTIAL

Bidding, Dynamic Allocation (DA), Enhanced Dynamic Allocation (EDA), tying DFP to AdX, and Privacy Sandbox programs and practices entailed one or more of the following: (a) failures to adequately or timely disclose changes to the auction's mechanics and purposes; (b) unwarranted restrictions on material information needed by auction participants and intermediaries; (c) denials of equal and fair access to inventory, demand, and functionality to advertisers, publishers, ad servers, exchanges, or ad buying tools; and (d) conflicts of interest. Those programs and practices jeopardized, and detrimentally affected, transparency and fairness of the auctions in which they were employed.

B. Introduction

342. I have read the descriptions and analyses in the Expert Report of Dr. Matthew Weinberg (the "Weinberg Report") with respect to certain of Google's conduct, including: (a) Dynamic Allocation (DA) and Enhanced Dynamic Allocation (EDA); (b) Header Bidding; (c) Unified Pricing Rules (UPR); (d) the two version of Dynamic Revenue Sharing (DRS); (e) Project Bernanke and its related versions; and (f) Reserve Price Optimization (RPO). Relevant excerpts from the Weinberg Report are attached hereto as Appendix D. I also have knowledge and information Google programs and practices as a result of my review of case materials (depositions, production documents, and pleadings) and, in some cases, my personal experience in the ad tech industry and the digital advertising ecosystem. Based on the facts from the Weinberg Report set forth in Appendix D, as well as my own knowledge and understanding, I have formulated certain opinions about Google's conduct, programs, and practices. That conduct generally falls into four categories:

- 1) Failures to disclose changes to auction rules.
- 2) Failures to provide equal access and functionality to competitors and customers.
- 3) Restricting the flow of information.
- 4) Conflicts of interest and self-dealing.

343. I address each of those categories below in subsections C–G.

C. Google Failed to Disclose Changes to Auction Rules

344. Inadequate or untimely disclosures of changes to auction rules harm advertisers, publishers, and other ad tech competitors. To effectively participate in an open web display programmatic auction, participants need to know the auction rules. These rules include the mechanics of the auction and criteria for determining the winning bid, the price paid by the winning bidder, the proceeds to the publisher, the take rate fee or commission to auction intermediaries, and other information affecting participants' strategies and decisions.

HIGHLY CONFIDENTIAL

345. As mentioned above, I have participated in both first-price and second-price auctions on both the buy side and the sell side. Advertisers must adopt a much cannier approach to first-price auctions, using bidding strategies that are inefficient but allow discovery of competitive bids. Publishers also change their price floor strategies in response to the auction format.

346. Changes to auction rules must be disclosed accurately, fully, and in a timely fashion. Accurate disclosure means truthful and correct information without misrepresentation. Full disclosure means no material omissions. Timely disclosure means information is provided promptly, not after-the-fact. Effective disclosure means information is conveyed clearly and understood by participants.

347. Without accurate, full, and timely disclosure, advertisers and participants cannot optimize their bidding strategies or engage with the auction effectively. Advertisers and publishers expect transparent and fair auctions, and undisclosed rule changes violate these expectations. Several of Google's programs involved failing to disclose changes to auction rules, affecting participant behavior and auction outcomes.

348. As stated, publishers and advertisers use different strategies for first-price and second-price auctions. If they believe an auction is second price when it is first price, their decision-making is skewed, impacting their results. With Reserve Price Optimization (RPO), advertisers believed they were in a standard second-price auction, but Google set artificially high reserve prices to optimize AdX revenue, often higher than the publisher's floor price. By using RPO, Google used advertisers' bids against them in future auctions, a fact that was reported on by DigiDay,¹⁶⁵ [REDACTED] and obfuscated by Google.¹⁶⁷

349. In Project Bernanke and a later iteration, Global Bernanke, Google manipulated auctions to increase how often the Google Display Network (GDN) won. For my understanding of this conduct and its basic steps, I partially rely on the expert report of Matthew Weinberg.¹⁶⁸

350. Project Bernanke began with auctions run on GDN. Google then took the top two bids in those auctions and altered them before sending them to AdX. The following diagram helps to illustrate this process:¹⁶⁹


¹⁶⁵ Digiday. "Google Sweetens Deal for Publishers With Dynamic Price Floors" (March 5, 2015). Accessed on June 4, 2024. <https://digiday.com/media/google-sweetends-deal-publishers-dynamic-price-floors/>.

¹⁶⁶ [REDACTED]

¹⁶⁷ GOOG-AT-MDL-003164171 at -172. "Re: [adx-questions:18269] Re: [programmatic-news] re: Google sweetens deal for publishers with dynamic price floors via Digiday" (March 6, 2015). Internal newsletter on programmatic news.

¹⁶⁸ See Appendix D.

¹⁶⁹ GOOG-NE-11753797 at -837.

HIGHLY CONFIDENTIAL

351. Let the top two bids be, in order, b_1 and b_2 , as in the diagram. The CAT2 auction is the auction within GDN. As shown by the vertical arrow, those bids received multipliers before being passed into AdX. The multiplier on b_1 , α , ranged from [REDACTED] while β ranged from [REDACTED]. The impact of these multipliers is that, in some auctions, GDN helped publishers at the expense of advertisers. And in some auctions, via β , GDN helped advertisers at the expense of publishers. Professor Weinberg has a full description of the program.

352. The primary impact, however, is that these multipliers were set up to maximize Google's revenue, as the slide indicates. This revenue maximization stemmed from the increased number of auctions in AdX that GDN would win.

353. A later iteration of Bernanke, called Global Bernanke continued to target an average take rate of [REDACTED] but set a [REDACTED] floor for the minimum take rate applied to impressions of any individual publisher, and a floor of [REDACTED] on the revenue an individual publisher would make through AdX.¹⁷⁰ Project Global Bernanke maintained a single pool for all publishers across AdX instead of individual per-publisher pools.

354. Both Bernanke and Global Bernanke created an environment harmful to publishers, advertisers, and other exchanges. While each group might earn, lose, or remain neutral with respect to profit, each group was harmed by the deceptive manipulation.

355. Publishers were harmed because the auction process often led to publishers showing lower quality ads than ads placed through other auctions. It is likely that publishers would have opted out of many auctions had they known that the rules had changed.

¹⁷⁰ GOOG-DOJ-AT-02471194, p. 1.

HIGHLY CONFIDENTIAL

356. Advertisers were harmed, especially those not on the GDN network, because the win rate was higher for GDN advertisers than non-GDN advertisers. GDN advertisers won more of the available impressions than they would without Bernanke and Global Bernanke.

357. Other exchanges were harmed because AdX completed manipulated auctions that would not have been completed otherwise, and therefore would have been available to other auctions.

358. After the iterations of Bernanke and Global Bernanke, Google adopted First-Price Bernanke, or first price auctions. The complex implementation is detailed in Professor Weinberg's report. As he writes, "The general framework of colluding and overbidding still apply, but the precise mechanics differ."¹⁷¹

359. In Project Bernanke, participants believed they were in a second-price auction, but it was essentially a third-price auction, with the publisher receiving the third-highest bid, the advertiser paying the second-highest bid, and Google pooling the difference to manipulate other auctions. The third-price auction harmed publishers, who may have changed their business decisions if they were aware of the practice. This undisclosed pooling prevented participants from making optimal decisions and harmed advertisers and other exchanges.

360. RPO, DRSp1,¹⁷² and Bernanke each represented a change to the auction rules that auction participants and competing exchanges would have expected to govern their auctions. For some period of time, Google failed to disclose these changes to participants in reasonable ways. It is my opinion that those undisclosed Google rule changes were contrary to the expectations of the auction participants and made it impossible for auction participants and competing exchanges to understand the rules that governed and applied to auctions run by Google, skewing decision-making and outcomes.

D. Google Failed to Provide Equal Access and Functionality to Competitors and Customers

361. Unequal access to information in an open web display auction creates imbalances that disadvantage some participants while advantaging others. This can distort auction outcomes and tip the scales in favor of the informationally advantaged participants.

362. Providing unequal access to information material to open web display campaigns and bidding strategies risks ineffective decision-making and

¹⁷¹ Appendix D at paragraph 264.

¹⁷² Further it is my understanding that Prof. Weinberg opines that disclosures around other versions of DRS were incomplete and misleading.

HIGHLY CONFIDENTIAL

discriminatory advantages. Control over certain functionalities in the ad tech ecosystem leads to self-preferencing, favoritism, and bias.

363. The detrimental consequences of unequal access are exacerbated when not properly disclosed by those controlling the information or functionality.

364. Several of Google's programs involved denying equal access to information or functionality. This unequal access violated participant expectations and negatively affected auction outcomes and campaign effectiveness.

365. For instance, Google's Dynamic Allocation (DA) gave AdX an exclusive "right of first refusal" on bids submitted on rival exchanges, allowing Google to submit winning bids with confidence. This unequal access affected participant decision-making and auction outcomes.

366. With Exchange Bidding, Google impaired rival exchanges' participation with AdX, limited publishers' ability to use header bidding auctions, and impeded performance measurement of header bidding. This inhibited header bidding use, affecting auction outcomes.

367. Google also tied AdX access to using DFP, foreclosing publishers from selling through AdX without DFP, denying critical functionality and interoperability.

368. The Privacy Sandbox initiative will replace third-party cookies with Google-specific API data, disadvantaging advertisers using competing tools. Advertisers will be disadvantaged by not being able to use the most accurate and sophisticated approaches to attribution modeling.

369. In summary, Google's programs like DA, EDA, and Exchange Bidding, and its tying of DFP to AdX denied participants equal access to log-level and bid-level information and functionality. The Privacy Sandbox will limit access to targeting information and post-campaign analysis, favoring Google's ad tech platforms.

E. Google Restricted the Flow of Information.

370. Possession and control of auction information give rise to potential dangers when access is restricted. Blocking or limiting material information results in less effective strategies and enables self-preferencing by the controlling party.

371. Several of Google's programs involved restricting information, contrary to participant expectations, affecting auction outcomes and campaign effectiveness.

372. Under Unified Pricing Rules (UPR), Google denied publishers the ability to set different reserves for different exchanges and tools, benefitting AdX.

HIGHLY CONFIDENTIAL

373. The Privacy Sandbox will eliminate third-party cookies, previously used for bidding and campaign efficacy, limiting new API targeting information to Google's platforms.

F. Google Exploited Conflicts of Interest and Engaged in Self-dealing

374. When a company operates a publisher ad server, ad buying tools, and an exchange, conflicts of interest arise. The potential for conflict is greater when the party's presence limits participants' options for comparable services. Conflicts arise when the party with control of material information or functionality discriminates among participants, violating expectations of fairness and impartiality.

375. In a competitive market, conflicts of interest can be mitigated by the presence of multiple service providers, allowing participants to switch if they perceive any unfair treatment. When a dominant player controls key aspects of the ad tech ecosystem—such as a publisher ad server, ad buying tools, and an exchange—and restricts viable alternatives, however, the potential for exploiting these conflicts increases. Without competitive options, the dominant party can leverage its position to favor its own interests, ultimately harming consumers and stifling competition.

376. Several of Google's programs involved conflicts of interest, favoritism, and self-preferencing, contrary to participant expectations, affecting auction outcomes.

377. Undisclosed conduct benefiting publisher-customers or Google at the expense of advertiser-customers, and vice versa, is conflict-ridden. Conduct benefiting Google's customers at the expense of competitors is conflicted when fairness is expected.

378. Dynamic Revenue Sharing (DRS) and DRSv1 led participants to believe they were in a second-price auction with consistent AdX take rates, but Google manipulated take rates to win impressions it would not have otherwise won. DRSv2 expanded this manipulation to raise take rates to compensate for lost fees. These changes, without disclosure, violated participant expectations and impacted auction outcomes.

379. The Weinberg Report identifies benefits and disadvantages of Google's programs like DA, EDA, Exchange Bidding, UPR, DRS, Bernanke, and RPO. These involved self-preferencing or conflicts of interest, affecting decision-making and auction outcomes. Google's tying of DFP to AdX and the Privacy Sandbox initiative are also acts of self-preferencing, reflecting conflicts of interest.

HIGHLY CONFIDENTIAL

G. Google Engaged in Experiments that Exploited its Users

380. Each of EDA, DRS, Bernanke, RPO, and UPR originated as “yield optimization experiments.”¹⁷³ Google’s yield optimization experiments are experimental algorithms applied to actual, live auctions to test whether those algorithms will increase Google’s revenue.¹⁷⁴ Google has performed experiments on auctions since the 2012-2013 timeframe.¹⁷⁵ Google usually conducts experiments on [REDACTED] of live auctions.¹⁷⁶ Google will run experiments on [REDACTED] or more of live auctions at times.¹⁷⁷ If the experiment successfully increases Google’s revenue in a way that outweighs its perceived negative impacts, then Google will “launch” the algorithm so that it is applied to all auctions.¹⁷⁸

381. As I discuss elsewhere in my report, Google’s conduct at issue in this case altered auction behavior such that, unknown to its users, Google auctions did not operate as second price auctions. As I also discuss, auction participants on the buy side have specific expectations about how second-price auctions function, specifically that buyers can bid the true value of the inventory to them without penalty. Google was fully aware that experiments may cause auctions not to behave like second price auctions. For example, Google made changes to its help center documentation in 2014 to remove “second price” from the description of its “Open” and “Private” auctions.¹⁷⁹ A change from a second-price auction is a big change to auction dynamics, but removing language from a help center page is not a highly visible method of disclosure. Prior to the 2014 change, Google’s then “current documentation on the behaviour of the AdX second price auction restrict[ed] the ability to experiment with auction dynamics and min CPM settings.”¹⁸⁰ Google’s documentation explains that to “maximiz[e] revenue and/or fill rate we need the right to not be restricted to a pure 2nd price auction.”¹⁸¹ Google chose to change its documentation rather than disclose

¹⁷³ [REDACTED]. 90:22-91:19 (agreeing that EDA, DRS, RPO, and UPR are yield optimization experiments), 98:10-103:13 (EDA and DRS were implemented on live auction traffic), 111:2-9 (UPR and RPO were implemented on live auction traffic); [REDACTED] 242:17-244:2 (testifying that Bernanke was an experiment applied to live auctions).

¹⁷⁴ [REDACTED] 94:14-22 (These four yield optimization experiments increased “both publisher and Google’s revenue”), 95:5-9 [REDACTED] could not think of any yield optimization ideas that did not increase Google’s revenue).

¹⁷⁵ Nitish Korula Dep., 412:9-15.

¹⁷⁶ [REDACTED] 106:14-23.

¹⁷⁷ Nitish Korula Dep., 491:3-17 (when asked about the “high end” of live traffic, Korula testified that it was not common “to go above [REDACTED]”), 413:2-414:11 (when Google gained more confidence in an experiment, increased percentages of live traffic to [REDACTED]).

¹⁷⁸ *Id.* at 413:2-414:11 (when Google gained more confidence in an experiment, it increased percentages of live traffic to [REDACTED], then launched it, and it was at [REDACTED]); [REDACTED] 136:8-14 (turning on an experiment for all of a publisher’s traffic is called “launching” it).

¹⁷⁹ GOOG-AT-MDL-016487180; [REDACTED]. 173:11-24 [REDACTED] agrees that the redlines for the help center documentation remove “second price” from the auction description and replaces it with “based on the highest net bid submitted”).

¹⁸⁰ GOOG-AT-MDL-016487180.

¹⁸¹ *Id.*

HIGHLY CONFIDENTIAL

its experiments to customers and seek their permission via the ability to opt in because “new opt in options for publishers and buyers will give us some flexibility but not large coverage for testing”¹⁸² and to avoid changing contracts with its users “which minimizes implementation cost and time.”¹⁸³

382. Google’s implementation of experiments that made AdX a non-second-price auction is surprising because in 2014 AdX was “a pure 2nd price auction” and Google referred to that as a “competitive advantage.”¹⁸⁴ Google noticed that “auction behavior is not defined in . . . buyer . . . or seller contracts” and that “[a]ll references to 2nd price auctions are contained with AdX buyer and AdX seller help center documentation.”¹⁸⁵ In short, Google determined that changing its help center documentation facilitated altering auction behavior without customer consent or re-negotiating contracts.

383. Google’s help center replaced the second price auction terminology with a convoluted description of auctions where “DoubleClick Ad exchange determines the winning bidder.”¹⁸⁶ Notably AdX’s “determination” is “**based on** the highest net bid submitted.”¹⁸⁷ The documentation does not say that the highest net bid wins, but merely that the result in some way is based on the highest net bid.¹⁸⁸ This statement is the creation of a technical loophole, not disclosure with an aim of educating customers. Further, Google’s documentation does not define the “highest net bid.” Google’s employee also could not clearly define this term at deposition and testified that it depends on the context as it is a term Google itself uses differently.¹⁸⁹ The help center documentation merely states that Google “may run limited experiments designed to optimize the auction.”¹⁹⁰ The documentation also does not say that “optimize” means increase Google’s revenue.¹⁹¹ Google’s documentation does not disclose that some experiments fail to optimize revenue and may result in lower revenue than the same auction would without the experiment.¹⁹² Similarly, the documentation does

¹⁸² *Id.*

¹⁸³ *Id.*

¹⁸⁴ *Id.* at 7183.

¹⁸⁵ *Id.*; see also *id.* (Listing the “Need to remove barriers to experimentation,” Opt-ins only get us part way there,” “Transparency is important,” and the “need to learn what works”).

¹⁸⁶ *Id.* at 7185.

¹⁸⁷ *Id.* at 7185 (emphasis added).

¹⁸⁸ *Id.*

¹⁸⁹ ██████████. 182:4-186:17 ██████████ testifying that “net bid” is the bid price after Google’s revenue share is taken out and also that “the exact meaning of net bid [is] determined by the context” and that he “can’t comment” on every one of these contexts).

¹⁹⁰ GOOG-AT-MDL-016487185.

¹⁹¹ See ██████████. 94:14-22 (These four yield optimization experiments were designed to increase “both publisher and Google’s revenue”), 95:5-9 ██████████ could not think of any yield optimization ideas that did not increase Google’s revenue). ██████████

¹⁹² See ██████████. 46:18-51:12 ██████████ testifying that in some experiments of DRS, publishers “receive[d] less than ██████████ revenue share in some transactions,” and that these publishers were not told or reimbursed).

HIGHLY CONFIDENTIAL

not disclose that, for example, publisher-side experiments would increase publisher revenue at the expense of buyers.¹⁹³

384. Google's witnesses confirmed that Google users were not informed about the specifics of experiments or their impact. For example, [REDACTED] testified that, in the hundreds of Google contracts [REDACTED] has reviewed, Google has not disclosed anything to publishers with respect to experiments that Google may or will run as to those publishers.¹⁹⁴ Meanwhile, [REDACTED] testified that customers are not informed when experiments are run on auctions that affect the win rates of impressions.¹⁹⁵ [REDACTED] further testified that customers do not have the option to opt out of experiments generally, which is why they are not informed when experiments are run that affect them.¹⁹⁶

385. Google also removed internal barriers and checks so that its engineers could freely conduct experiments. [REDACTED]

[REDACTED] explained in [REDACTED] deposition that the process for designing and implementing an experiment takes as little as one week.¹⁹⁷ An engineer may propose an experiment, and so long as it shows the potential to increase Google's revenue, a Google engineering team lead may approve the experiment for live auctions.¹⁹⁸ While some experiments received additional review from Google's legal department, escalating an experiment for additional scrutiny depends on the discretion of Google's engineers.¹⁹⁹ When an engineer requests approval to run an experiment, the engineer could do so in as informal a manner as sending an email requesting approval to run an "unconstrained" experiment on live traffic.²⁰⁰ Google's 2014 documentation change effectively streamlined the process for Google to experiment on auctions with new algorithms to increase its revenue with little internal oversight and without disclosing the details of its experiments to users.

¹⁹³ [REDACTED]. 94:14-22 (admitting that EDA, DRS, UPR, and RPO all increase publisher's revenue and Google's revenue); 211:21-212:10 (agreeing that the "point" of RPO "is for advertisers or buyers to pay more in an auction so that the publishers receive more revenue"); 334:20-335:8 ("Q. Even some of Google's successful experiments, those that maybe increase publisher revenue, might be harmful, financially, to advertisers, right? ...Causing them to pay more money than they should be paying.... THE WITNESS: I think some experiments or some yield optimization increase publisher revenue, and some buyer need to pay more, yes."); [REDACTED] 167:15-168:1 (testifying that RPO could make queries more expensive for buyers); GOOG-AT-MDL-B-004637455 (RPO launch document explaining that "Setting optimized prices on behalf of publishers makes queries more expensive for buyers").

¹⁹⁴ [REDACTED]. 295:16-296:9.

¹⁹⁵ [REDACTED]. 305:2-311:22.

¹⁹⁶ *Id.* at 319:3-19.

¹⁹⁷ [REDACTED] 159:23-160:18.

¹⁹⁸ *Id.* at 113:9-114:8, 116:4-15.

¹⁹⁹ *Id.* at 115:1-116:15.

²⁰⁰ [REDACTED] 249:16-255:23 ([REDACTED] agreeing he sent an email requesting approval to run an "unconstrained" experiment on Bernanke because Bernanke was not profitable but could not remember what he meant by "unconstrained"); *see also* GOOG-AT-MDL-002762069.

HIGHLY CONFIDENTIAL

386. Google also allowed its engineering teams broad autonomy over the auctions that are subject to experiments. [REDACTED] testified that hundreds of experiments may be running each day and Google does not track what percentage of auctions are subject to experiments.²⁰¹ Indeed Google's engineers believed it was best that experiments randomly affect auctions.²⁰² Some auctions were also subject to multiple experiments at once.²⁰³

387. Indeed, Nitish Korula, Google's corporate representative, testified that as of August 2019 there were [REDACTED] experiments across sell and buy-side display ads,²⁰⁴ and further explained a single auction (query) would, on average, have some [REDACTED] experiments applied to it.²⁰⁵ In a Google blog post from 2016 describing RPO, Google declared that "we have applied optimized pricing to about 15% of transactions creating over 5% lift in revenue for publishers using the Open Auction."²⁰⁶ Applying an experiment to [REDACTED] of auctions impacts a substantial number of auctions because Google runs billions of auctions each day.²⁰⁷ Google documents and Nitish Korula's testimony demonstrate that, at least as of 2019, Google was conducting some [REDACTED] auctions per day.²⁰⁸ But, even assuming Google was only conducting 1 billion auctions per day, an experiment running on 1% of traffic would be applied to 10 million auctions.²⁰⁹ At an average clearing price of \$2.50, a single experiment applied to 1% of auctions would concern \$25 million in auctions.²¹⁰

388. Google's experiments may also run on traffic for extended periods of time. For example, DRS v1 began with experiments on live traffic that started in late 2014.²¹¹ Google continued to experiment on live auctions until DRS v1 was fully launched in summer of 2015.²¹² Experiments related to DRS v2 began by no later than January 2016.²¹³ The unmonitored nature of experiments is further exemplified by [REDACTED] testimony that experiment codes have an expiration date, because otherwise "some experiment" could "just go[] on forever and people forget about it."²¹⁴

²⁰¹ [REDACTED]. 269:12-270:12.

²⁰² *Id.* at 118:14-22 ("[W]e believe that random traffic sampling will give us and help us to [make the] best decision for publisher[s] on yield optimization.").

²⁰³ *Id.* at 122:3-11 (agreeing that "[i]t's possible that the auctions that are affected by these experiments are impacted by more than one experiment at the same time").

²⁰⁴ GOOG-AT-MDL-B-001646464 at 6468; *see also* Nitish Korula Dep. 174:16-175:13.

²⁰⁵ GOOG-AT-MDL-B-001646464 at 6468; Nitish Korula Dep. 178:4-21.

²⁰⁶ Nitish Korula Dep. Ex. 304.

²⁰⁷ [REDACTED] 107:4-16.

²⁰⁸ GOOG-AT-MDL-B-001646464 at 6466.

²⁰⁹ *Id.* at 107:20-24 ([REDACTED] agreeing that an experiment running on 1% of 1 billion auctions per day would affect 10 million auctions per day).

²¹⁰ *Id.* at 110:4-20.

²¹¹ *Id.* at 223:2-14.

²¹² *Id.* at 229:22-23 (agreeing DRS would have been launched by summer 2015).

²¹³ *Id.* at 233:10-18.

²¹⁴ [REDACTED]. 285:2-8.

HIGHLY CONFIDENTIAL

389. Google's experiments often have bugs that cause problems, and Google does not reimburse or inform customers when that happens. For example, during RPO experiments, Google recognized that RPO "hurts publishers where their floors are already high."²¹⁵

390. When asked about an email exchange between [REDACTED] Nitish Korula, [REDACTED] explained that Google tried to make sure it is not "not hurting publishers too much during development" to ensure "a win for our publishers."²¹⁶ But Google documents indisputably confirm instances where experiments harmed its publishers. For instance, a Project Bernanke experiment resulted in "publisher payout drop[ping] by [REDACTED]."²¹⁷

391. Later, Google also experienced issues with RPO and DRS experiments overlapping on certain auctions, causing Google to lose revenue.²¹⁸ Google took active steps to conceal its experiments because it knew that they were duplicitous in messaging to publishers and advertisers. For instance, an email thread between Google personnel from 2014 makes clear that Google's response to publisher concerns about auction behavior was to avoid announcing experiments and to hide reserve prices that publishers would pay.²¹⁹

392. In another instance, Google employees favored excluding certain publishers with bid data access from an experiment because they were able to detect "anomalies" and would "notice experiment traffic in their bid data files."²²⁰ Similarly, Google took pains to roll out RPO during April because traffic was more volatile, making it harder for customers to detect the change: "Ramp up to 100% in the first (2?) weeks of April, when fluctuations are common anyway."²²¹

393. And Google created a "sensitive publishers list" for purposes of experiments.²²² The list consisted of large, sophisticated publishers—ones better equipped

²¹⁵ *Id.* at 182:14-183:10; GOOG-AT-MDL-018794834.

²¹⁶ [REDACTED] 183:11-184:17; GOOG-AT-MDL-018794833-35.

²¹⁷ GOOG-AT-MDL-012692665; *see also* GOOG-AT-MDL-014566659 (instance where "overall rtb revenue lift due to RPO was actually negative").

²¹⁸ GOOG-DOJ-14008627 at 8631; *see also* [REDACTED] 291:11-294:19 ([REDACTED] testifying that Google internal document shows that RPO and DRS experiments caused a 12-hour outage on August 25, and on August 27th, after Google engineers went "all hands on deck," they "get to the bottom" of the problem.); *see also* GOOG-AT-MDL-014566659-60 (email from [REDACTED] that the RPO/DRS outage and resulting revenue loss "can't be attributed to a handful of buyers or pubs . . . and seems spread out across pubs and buyers"); GOOG-DOJ-27760516 (explaining that, in response to experiment risk, Google *would not* take the precaution of "Regular refunding of harm done to pubs and/or buyers").

²¹⁹ *See* GOOG-DOJ-32277200.

²²⁰ GOOG-AT-MDL-008962081.

²²¹ GOOG-DOJ-13997420.

²²² GOOG-DOJ-AT-02191375; [REDACTED] 275:16-24; 279:6-280:7.

HIGHLY CONFIDENTIAL

to detect Google's experiments.²²³ And Google excluded publishers from experiments if they were "litigious."²²⁴

394. In an email thread, Nitish Korula suggests that some of Google's problems "go away" if Google does not "actually announce reserve prices."²²⁵ In response, [REDACTED] suggests "experimenting with selectively obscuring the reserve price we send to buyers," because Google is "potentially doing things that don't quite fall under the description of 'truthful second price auction.'"²²⁶

395. Google understood the risks associated with its experiments. [REDACTED] in an email to the team designing yield optimization experiments that included [REDACTED] Nitish Korula, wrote about the risk of Google's experimentation on its relationship with its partners.²²⁷ [REDACTED] specifically acknowledged that experiments created "risk towards [Google's] partners."²²⁸ Nevertheless, [REDACTED] wrote that such risk was worthwhile and could be mitigated by "excluding sensitive partners (pubs and buyers) and keep the traffic fraction low, but then just accept and deal with any left-over risk – e.g. not require partners to opt-in at an experiments stage."²²⁹ Other Google internal documents support this, including GOOG-DOJ-27760516, which instructs that Google might take the precaution of "Exclusions of some pubs [REDACTED]" in response to risks carried by experiments.²³⁰

396. Google witnesses testified that any changes to its auctions under the veil of an "experiment" (1) could harm users and (2) required no disclosure to users. For example, [REDACTED] testified that "when a publisher receives less money for a given impression, that could be perceived as harm."²³¹ [REDACTED] agreed that "[i]t's likely that a customer would want to know if an experiment is not benefitting them."²³² Nonetheless, [REDACTED] could not think of any example where he explained to customers what type of experiments Google ran.²³³ Google's other witnesses similarly could not recall any instance where Google reimbursed its customers where experiments skewed results in a way that either underpaid publishers or overcharged advertisers.²³⁴ Further, [REDACTED] could not think of any example where a customer was harmed by an experiment and that customer and Google had a conversation about it.²³⁵ Google internal documents support this approach. For instance, GOOG-

²²³ See [REDACTED] 284:8-11.

²²⁴ GOOG-AT-MDL-015179479.

²²⁵ GOOG-DOJ-32277200 (June 20, 2014, e-mail at 9:51 am from Nitish Korula).

²²⁶ *Id.* (June 20, 2014, e-mail at 10:11 am from [REDACTED]).

²²⁷ GOOG-AT-MDL-019653416.

²²⁸ *Id.*

²²⁹ *Id.*

²³⁰ See GOOG-DOJ-27760516.

²³¹ [REDACTED] 138:14-21.

²³² *Id.* at 136:19-137:1.

²³³ *Id.* at 118:10-22.

²³⁴ Nitish Korula Dep. 179:23-180:7.

²³⁵ *Id.* at 118:23-119:4.

HIGHLY CONFIDENTIAL

DOJ-27760516 indicates that precautions that Google would *not* take in response to experiments that carry risk include “Communications around individual experiments” and “Publisher opt-in to experiments.”²³⁶ Nitish Korula’s testimony makes clear that the way Google structured its experiments, Google tracked the generalized trends in revenue caused by its experiments in the form of either positive or negative “revenue lift.”²³⁷

XI. Reservation of Rights

397. My opinions and analysis are based upon the information available to me to date. I may review and consider additional information that may be produced by the parties to this dispute. I intend to supplement my opinions based upon that review, if it is appropriate to do so. I also reserve the ability to provide rebuttal opinions and testimony in this matter, to create demonstratives for use at trial based upon the information contained in this report, appendices, and exhibits, and generally to utilize other graphical depictions as aids in the presentation of my findings.

398. This report is confidential. It is only to be used for its stated purpose and is subject to a protective order.

²³⁶ GOOG-DOJ-27760500 at 0516.

²³⁷ See generally, Nitish Korula Dep. 158:7-201:6 (describing analysis of experiments as statistical trends compared to expectations to determine where experiments caused either a positive or negative revenue lift and also being unable to quantify harm to any individual Google publisher or advertiser).

APPENDIX A

JOHN CHANDLER

2301 Milwaukee Ave, Minneapolis, MN 55404
+1.406.544.8720
john.chandler@gmail.com, john.chandler@business.umt.edu, john@datainsightsllc.com

EDUCATION

University of Montana
Ph.D. in Mathematical Sciences **2010**
Dissertation: Statistical Learning in Online Advertising

University of Washington
MS in Mathematics **1999**
Thesis: Visualizing Network Flows

Middlebury College
BA in Mathematics (Magna Cum Laude, Thesis Honors) **1996**
Minor: Sociology
Thesis: Representations of Finite Groups

AWARDS

MS in Business Analytics Outstanding Faculty Award **2017-2022**
Outstanding Non-Tenure Track Faculty Award **2017, 2020**

TEACHING EXPERIENCE

University of Montana
Clinical Professor **2014-Present**
Co-developed curriculum for MS in Business Analytics. Developed courses Applied Data Analytics, Text Mining and Unstructured Data, Advanced Applied Modeling, Introduction to SQL, and Telling Stories with Data. Additionally, taught weekend and short courses.

Instructor **2007**
Developed syllabus and taught “Simulations in R”.

ORT University, Montevideo Uruguay
Visiting Professor **2018-Present**
Annual teaching of Telling Stories with Data for analytics students and assistance with the development of business analytics degree.

University of San Diego
Adjunct Professor **2021-Present**

Teaching Applied Text Mining (ADS-509) in the MS of Applied Data Science degree.

University of Washington

Instructor

1998-1999

Taught Calculus, Pre-Calculus, Linear Algebra, and Algebra

RELATED EXPERIENCE

Data Insights

Managing Partner

2012 – Present

Managed data science consulting practice, working with numerous Fortune 500 companies and start-ups.

Microsoft

Research Director

2007-2011

Data scientist at Microsoft, working across a variety of advertiser, publisher and auction products. Last position was leading research for Microsoft TV.

aQuantive

Principal Analyst

1999-2007

Data scientist working at the intersection of product, statistics and marketing.

BOOKS

Steele, B, J Chandler and S Reddy (2018). *Algorithms for Data Science*. New York: Springer-Verlag

PUBLICATIONS AND PAPERS

Smith, M. L., MacLehose, R. F., Chandler, J. W., & Berman, J. D. (2022). Thunderstorms, Pollen, and Severe Asthma in a Midwestern, USA, Urban Environment, 2007–2018. *Epidemiology*, 33(5), 624-632.

Earnest, D., & Chandler, J. (2021). Making time: Words, narratives, and clocks in elementary mathematics. *Journal for Research in Mathematics Education*, 52(4), 407-443.

Yung, L., Chandler, J., & Haverhals, M. (2015). Effective weed management, collective action, and landownership change in western Montana. *Invasive plant science and management*, 8(2), 193-202.

Yung, Laurie., Freimund, Wayn., & Chandler-Pepelnjak, John. (2008). Wilderness politics in the American West. *International Journal of Wilderness*, 14(2), 14-23.

WORKING PAPERS

Chandler, J. & Bu, X. "We will but I did: Collectivism versus individualism in political convention speeches". *Preparing for initial submission in PNAS*

Metcalf, A., Birdsong, M., & Chandler, J. "Private lands in the public trust". *Preparing for initial submission to Society and Natural Resources*

Chandler, J. "Violations of independence in network statistics". *Preparing for initial submission at Social Network Analysis*

CONSULTING REPORTS AND WHITE PAPERS

"Measuring ROI Beyond the Last Ad"	2009
"The Long Road to Conversion: The Digital Purchase Funnel"	2008
"Optimal Frequency: The Impact of Frequency on Conversion Rates"	2006
"Traditional Advertising Metrics on the Web: Forecasting GRP's, Reach and Effective Reach Online"	2007
"Forecasting Reach, Frequency, and GRPs on the Internet"	2004
"Atlas Annual Holiday Shopping Report"	2000-2007

PATENT APPLICATIONS

"System and method for determining internet advertising strategy"	
US20030074252A1; with A Easterly	2003

INVITED TALKS

"Teaching Analytics in the Age of AI"	
Universidad ORT Uruguay	2024
"DevOps is Data Science"	
Data Tech	2019
"Measuring Marketing Effectiveness"	
Universidad ORT Uruguay	2019
"Data Science is Dev-Ops"	
University of Montana	2018
"Marketing Science at Tableau"	
Tableau Conference	2018
"The Collision of Data Science and Marketing"	
University of Montana	2015
"Causal Inference is Hard", Panel Discussion	

Advertising Research Foundation 2014

"Understanding Attribution"

Advertising Research Foundation 2012

DEGREE COMMITTEES

- Tina Cummins, MS in Economics, 2017, *Impact of local government revenue and spending during oil and gas booms in the Rocky Mountain States*
- Omid Khormali, Ph.D. in Mathematica Sciences, 2019, *Extremal problems for forests in graphs and hypergraphs*
- Hannah Leonard, MS in Forestry, 2020, *A Conservation Marketing Toolkit: Systematic Literature Mapping, Microtargeting Conservation Easements, and Conservation Corridor Prioritization*
- Nate Bender, MS in Forestry, 2022, *Call your elected officials: Identifying predictors and audiences for collective climate action*
- Madeline Damon, MS in Forestry, 2023, *INVESTIGATING INTELLECTUAL DIVERSITY: A CRITICAL EXAMINATION OF ACADEMIC PUBLISHING PRACTICES AND THEIR EFFECTS ON WILDLIFE CONSERVATION*
- Tina Cummins, Forestry Ph.D., in progress
- Anna Marbut, Interdisciplinary Ph.D., in progress
- Chelle Twilliger, Ph.D in Forestry, in progress

LANGUAGES

English—native language

Spanish—speak, read, and write with competence

MEMBERSHIPS

American Marketing Association

American Statistical Association

APPENDIX B

HIGHLY CONFIDENTIAL

**APPENDIX B: PRIOR TESTIFYING
HISTORY OF DR. JOHN CHANDLER, Ph.D.**

1. In Re: JUUL Labs, Inc. Marketing, Sales Practices and Products Liability Litigation, Case No. 19-md-02913-WHO (N.D. Cal. 2019)
 - Deposed in July, 2021; October, 2021; and May, 2022.
2. The State of Alaska v. JUUL Labs, Inc., et al., Case No. 3AN-20-09477, In the Superior Court of Alaska, Third Judicial District, Anchorage, Alaska (2020)
 - Deposed in April, 2024.

APPENDIX C

HIGHLY CONFIDENTIAL

Appendix C: Materials Relied Upon & Materials Considered

MATERIALS RELIED UPON

1. All documents and sources referred to and cited in Dr. Chandler's report and its footnotes, including Bates-stamped documents, deposition transcripts, and other sources.

MATERIALS CONSIDERED

Discovery Responses

All available discovery responses produced within the matter of *The State of Texas, et al. v. Google*, Case Number: 4:20-cv-00957-SDJ, including:

1. The Parties' amended initial disclosures;
2. The Parties' discovery responses and objections to Interrogatories, Requests for Admission, and Requests for Production; and
3. Google's written responses to Plaintiffs' Rule 30(b)(6) Notice.

Deposition Transcripts & Exhibits

All available deposition transcripts and exhibits within the matter of *The State of Texas, et al. v. Google*, Case Number: 4:20-cv-00957-SDJ, including:



HIGHLY CONFIDENTIAL

HIGHLY CONFIDENTIAL

Category	Value
1	~10
2	0
3	0
4	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	~90

All available deposition transcripts and exhibits within the matter of *USA v. Google*, Case Number: 1:23-cv-00108-LMB-JFA, including:

ANSWER The answer is (A) $\frac{1}{2} \ln(1 + x^2)$.

HIGHLY CONFIDENTIAL



HIGHLY CONFIDENTIAL



Expert Reports & Declarations

All available expert reports (with redactions) within the matter of *USA v. Google*,
Case Number: 1:23-cv-00108-LMB-JFA, including:

1. Declarations of Google Employees
2. 2023.12.22 Expert Report of Gabriel Weintraub, GOOG-AT-MDL-C-000018734
3. 2023.12.22 Expert Report of R. Ravi, GOOG-AT-MDL-C-000019017
4. 2023.12.22 Expert Report of Robin S. Lee, GOOG-AT-MDL-C-000019273
5. 2023.12.22 Expert Report of Rosa Abrantes-Metz, GOOG-AT-MDL-C-000019786
6. 2023.12.22 Expert Report of Thomas S. Respass, GOOG-AT-MDL-C-000020106
7. 2023.12.22 Expert Report of Timothy Simcoe, GOOG-AT-MDL-C-000020274
8. 2024.01.13 Errata to Abrantes-Metz Expert Report, GOOG-AT-MDL-C-000020435
9. 2024.01.13 Errata to Ravi Expert Report, GOOG-AT-MDL-C-000020437
10. 2024.01.13 Errata to Respass Expert Report, GOOG-AT-MDL-C-000020440
11. 2024.01.13 Errata to Simcoe Expert Report, GOOG-AT-MDL-C-000020467
12. 2024.01.13 Errata to Weintraub Expert Report, GOOG-AT-MDL-C-000020471
13. 2024.01.23 Chevalier Expert Report, GOOG-AT-MDL-C-000020474
14. 2024.01.23 Ferrante Expert Report, GOOG-AT-MDL-C-000020714

HIGHLY CONFIDENTIAL

- 15.2024.01.23 Ghose Expert Report, GOOG-AT-MDL-C-000020767
- 16.2024.01.23 Israel Expert Report, GOOG-AT-MDL-C-000021036
- 17.2024.01.23 Milgrom Expert Report, GOOG-AT-MDL-C-000021794
- 18.2024.01.23 Rinard Expert Report, GOOG-AT-MDL-C-000022191
- 19.2024.01.23 Shirky Expert Report, GOOG-AT-MDL-C-000022229
- 20.2024.01.23 Simonson Expert Report, GOOG-AT-MDL-C-000022290
- 21.2024.01.23 Skinner Expert Report, GOOG-AT-MDL-C-000022948
- 22.2024.02.13 Expert Rebuttal Report of Adoria Lim, GOOG-AT-MDL-C-000023002
- 23.2024.02.13 Expert Rebuttal Report of Gabriel Weintraub, GOOG-AT-MDL-C-000023226
- 24.2024.02.13 Expert Rebuttal Report of Kenneth Wilbur, GOOG-AT-MDL-C-000023322
- 25.2024.02.13 Expert Rebuttal Report of R. Ravi, GOOG-AT-MDL-C-000023435
- 26.2024.02.13 Expert Rebuttal Report of Robin S. Lee, GOOG-AT-MDL-C-000023516
- 27.2024.02.13 Expert Rebuttal Report of Rosa Abrantes-Metz, GOOG-AT-MDL-C-000023887
- 28.2024.02.13 Expert Rebuttal Report of Timothy Simcoe, GOOG-AT-MDL-C-000024064
- 29.2024.02.13 Expert Rebuttal Report of Wayne Hoyer, GOOG-AT-MDL-C-000024138
- 30.2024.02.13 Expert Rebuttal Report of Wenke Lee, GOOG-AT-MDL-C-000024270
- 31.2024.02.16 Errata to Ravi Rebuttal Report, GOOG-AT-MDL-C-000024387
- 32.2024.02.20 Errata to Simcoe Rebuttal Report, GOOG-AT-MDL-C-000024389
- 33.2024.02.23 Errata to Weintraub Rebuttal Report, GOOG-AT-MDL-C-000024390
- 34.2024.02.23 Supplemental Errata to Weintraub Expert Report, GOOG-AT-MDL-C-000024391
- 35.2024.02.24 Errata to Wilbur Rebuttal Report, GOOG-AT-MDL-C-000024392
- 36.2024.02.26 Errata to Hoyer Rebuttal Report, GOOG-AT-MDL-C-000024397
- 37.2024.02.28 Errata to Abrantes-Metz Rebuttal Report, GOOG-AT-MDL-C-000024399
- 38.2024.03.04 Expert Supplemental Report of Robin S. Lee, GOOG-AT-MDL-C-000024403
- 39.2024.03.08 Consolidated Errata to Lee Rebuttal Report, GOOG-AT-MDL-C-000024436


HIGHLY CONFIDENTIAL

40. 2024.01.13 Expert Report of Weintraub Errata, GOOG-AT-MDL-C-000040965
41. 2024.01.13 Expert Report of Simcoe Errata, GOOG-AT-MDL-C-000040961
42. 2024.01.13 Expert Report of Respess Errata_with Figure Errata_Redacted, GOOG-AT-MDL-C-000040934
43. 2024.01.13 Expert Report of R Ravi Errata, GOOG-AT-MDL-C-000040931
44. 2024.01.13 Expert Report of Abrantes-Metz Errata, GOOG-AT-MDL-C-000040929
45. 2024.03.08 Consolidated Errata to Lee Rebuttal Report, GOOG-AT-MDL-C-000040926
46. 2024.03.04 Expert Supplemental Report of Robin S. Lee, PhD, GOOG-AT-MDL-C-000040893
47. 2024.02.28 Rebuttal Report Errata of Rosa Abrantes-Metz Signed, GOOG-AT-MDL-C-000040889
48. 2024.02.25 Expert Rebuttal Report of Hoyer Errata, GOOG-AT-MDL-C-000040887
49. 2024.02.24 Wilbur Rebuttal Errata, GOOG-AT-MDL-C-000040882
50. 2024.02.23 Weintraub Rebuttal Report Errata, GOOG-AT-MDL-C-000040881
51. 2024.02.23 Expert Report of Weintraub Supplemental Errata, GOOG-AT-MDL-C-000040880
52. 2024.02.20 Errata to Simcoe Rebuttal Report, GOOG-AT-MDL-C-000040879
53. 2024.02.16 Errata to Ravi Rebuttal Report (Highly Confidential), GOOG-AT-MDL-C-000040877
54. 2024.02.13 Rebuttal Report of Rosa Abrantes-Metz, GOOG-AT-MDL-C-000040700
55. 2024.02.13 Expert Report of Wenke Lee, GOOG-AT-MDL-C-000040583
56. 2024.02.13 Expert Rebuttal Report of Wayne Hoyer, GOOG-AT-MDL-C-000040451
57. 2024.02.13 Expert Rebuttal Report of Timothy Simcoe_Redacted, GOOG-AT-MDL-C-000040377
58. 2024.02.13 Expert Rebuttal Report of Robin S. Lee_Redacted, GOOG-AT-MDL-C-000040006
59. 2024.02.13 Expert Rebuttal Report of R Ravi, GOOG-AT-MDL-C-000039925
60. 2024.02.13 Expert Rebuttal Report of Kenneth Wilbur_Redacted, GOOG-AT-MDL-C-000039812
61. 2024.02.13 Expert Rebuttal Report of Gabriel Weintraub_Redacted, GOOG-AT-MDL-C-000039716
62. 2024.02.13 Expert Rebuttal Report of Adoria Lim_Redacted, GOOG-AT-MDL-C-000039492

HIGHLY CONFIDENTIAL

- 63.2024.01.23 Expert Report of William Clay Shirky, GOOG-AT-MDL-C-000039431
- 64.2024.01.23 Expert Report of Paul R. Milgrom, GOOG-AT-MDL-C-000039034
- 65.2024.01.23 Expert Report of Martin C. Rinard, GOOG-AT-MDL-C-000038996
- 66.2024.01.23 Expert Report of Mark A. Israel_Redacted, GOOG-AT-MDL-C-000038238
- 67.2024.01.23 Expert Report of Judith A. Chevalier_Redacted, GOOG-AT-MDL-C-000037998
- 68.2024.01.23 Expert Report of Itamar Simonson, GOOG-AT-MDL-C-000037340
- 69.2024.01.23 Expert Report of Douglas Skinner, GOOG-AT-MDL-C-000037286
- 70.2024.01.23 Expert Report of Anthony J. Ferrante, GOOG-AT-MDL-C-000037233
- 71.2024.01.23 Expert Report of Anindya Ghose_Redacted, GOOG-AT-MDL-C-000036954
- 72.2023.12.22 Expert Report of Timothy Simcoe_Redacted, GOOG-AT-MDL-C-000036793
- 73.2023.12.22 Expert Report of Thomas Respess_Redacted, GOOG-AT-MDL-C-000036625
- 74.2023.12.22 Expert Report of Rosa Abrantes-Metz_Redacted, GOOG-AT-MDL-C-000036305
- 75.2023.12.22 Expert Report of Robin S. Lee, PhD_Redacted, GOOG-AT-MDL-C-000035792
- 76.2023.12.22 Expert Report of R Ravi_Redacted, GOOG-AT-MDL-C-000035536
- 77.2023.12.22 Expert Report of Gabriel Weintraub_Redacted, GOOG-AT-MDL-C-000035253

Bates Stamped Productions, including:



HIGHLY CONFIDENTIAL

[REDACTED]	[REDACTED]	49.	GOOG-AT-MDL-003161451
		50.	GOOG-AT-MDL-004074544
		51.	GOOG-AT-MDL-004232880
		52.	GOOG-AT-MDL-004233138
		53.	GOOG-AT-MDL-004300268
		54.	GOOG-AT-MDL-004416785
		55.	GOOG-AT-MDL-004436768
		56.	GOOG-AT-MDL-004555181
		57.	GOOG-AT-MDL-006099844
		58.	GOOG-AT-MDL-006161050
		59.	GOOG-AT-MDL-006334729
		60.	GOOG-AT-MDL-006873424
		61.	GOOG-AT-MDL-006966530
40.	GOOG-AT-MDL-001263607	62.	GOOG-AT-MDL-007175167
41.	GOOG-AT-MDL-001390730	63.	GOOG-AT-MDL-007343585
42.	GOOG-AT-MDL-001391213	64.	GOOG-AT-MDL-007346556
43.	GOOG-AT-MDL-001933227	65.	GOOG-AT-MDL-007364833
44.	GOOG-AT-MDL-002105969	66.	GOOG-AT-MDL-007375672
45.	GOOG-AT-MDL-002105984	67.	GOOG-AT-MDL-007387750
46.	GOOG-AT-MDL-002124829	68.	GOOG-AT-MDL-007397182
47.	GOOG-AT-MDL-002390899	69.	GOOG-AT-MDL-007397197
48.	GOOG-AT-MDL-002393442		

HIGHLY CONFIDENTIAL

70.	GOOG-AT-MDL- 008148533 / GOOG-AT- MDL-008148529	91.	GOOG-AT-MDL- 012524006
71.	GOOG-AT-MDL- 008517788	92.	GOOG-AT-MDL- 012549335
72.	GOOG-AT-MDL- 008588684	93.	GOOG-AT-MDL- 012693796
73.	GOOG-AT-MDL- 008682082 / GOOG-AT- MDL-008682071	94.	GOOG-AT-MDL- 012767138
74.	GOOG-AT-MDL- 008754374	95.	GOOG-AT-MDL- 012837016
75.	GOOG-AT-MDL- 008835346	96.	GOOG-AT-MDL- 012857198
76.	GOOG-AT-MDL- 008858602	97.	GOOG-AT-MDL- 013290688
77.	GOOG-AT-MDL- 008881206	98.	GOOG-AT-MDL- 013291089
78.	GOOG-AT-MDL- 008886980	99.	GOOG-AT-MDL- 013292974
79.	GOOG-AT-MDL- 008953893	100.	GOOG-AT-MDL-013299524
80.	GOOG-AT-MDL- 008964888	101.	GOOG-AT-MDL-013299531
81.	GOOG-AT-MDL- 008979664	102.	GOOG-AT-MDL-013300202
82.	GOOG-AT-MDL- 008991390	103.	GOOG-AT-MDL-013378392
83.	GOOG-AT-MDL- 009026140	104.	GOOG-AT-MDL-013908958
84.	GOOG-AT-MDL- 009289718	105.	GOOG-AT-MDL-013918668
85.	GOOG-AT-MDL- 009291120	106.	GOOG-AT-MDL-014427012
86.	GOOG-AT-MDL- 009299907	107.	GOOG-AT-MDL-014460206
87.	GOOG-AT-MDL- 009321580	108.	GOOG-AT-MDL-014462378
88.	GOOG-AT-MDL- 009429957	109.	GOOG-AT-MDL-014486274
89.	GOOG-AT-MDL- 012512067	110.	GOOG-AT-MDL-014524447
90.	GOOG-AT-MDL- 012514705	111.	GOOG-AT-MDL-014618288
		112.	GOOG-AT-MDL-015241235
		113.	GOOG-AT-MDL-015622194
		114.	GOOG-AT-MDL-015844174
		115.	GOOG-AT-MDL-015929587
		116.	GOOG-AT-MDL-015997353
		117.	GOOG-AT-MDL-016457027
		118.	GOOG-AT-MDL-016534880
		119.	GOOG-AT-MDL-016627159
		120.	GOOG-AT-MDL-016656237
		121.	GOOG-AT-MDL-016772599
		122.	GOOG-AT-MDL-016838311
		123.	GOOG-AT-MDL-016924839
		124.	GOOG-AT-MDL-016937590
		125.	GOOG-AT-MDL-016943922

HIGHLY CONFIDENTIAL

126.	GOOG-AT-MDL-016967094	161.	GOOG-AT-MDL-B-002091565
127.	GOOG-AT-MDL-017187837	162.	GOOG-AT-MDL-B-002095353
128.	GOOG-AT-MDL-017394050	163.	GOOG-AT-MDL-B-002095501
129.	GOOG-AT-MDL-017494582	164.	GOOG-AT-MDL-B-002095769
130.	GOOG-AT-MDL-017664768	165.	GOOG-AT-MDL-B-002097533
131.	GOOG-AT-MDL-017746412	166.	GOOG-AT-MDL-B-002097570
132.	GOOG-AT-MDL-017749638	167.	GOOG-AT-MDL-B-002097648
133.	GOOG-AT-MDL-017762649	168.	GOOG-AT-MDL-B-002098265
134.	GOOG-AT-MDL-017864022	169.	GOOG-AT-MDL-B-002099366
135.	GOOG-AT-MDL-018248228	170.	GOOG-AT-MDL-B-002105135
136.	GOOG-AT-MDL-018427318	171.	GOOG-AT-MDL-B-002500395
137.	GOOG-AT-MDL-018448707	172.	GOOG-AT-MDL-B-002514153
138.	GOOG-AT-MDL-018548592	173.	GOOG-AT-MDL-B-002547489
139.	GOOG-AT-MDL-018618351	174.	GOOG-AT-MDL-B-002552122
140.	GOOG-AT-MDL-018652651	175.	GOOG-AT-MDL-B-002624643
141.	GOOG-AT-MDL-018998910	176.	GOOG-AT-MDL-B-002760309
142.	GOOG-AT-MDL-019001498	177.	GOOG-AT-MDL-B-002762758
143.	GOOG-AT-MDL-019306356	178.	GOOG-AT-MDL-B-002763194
144.	GOOG-AT-MDL-019386250	179.	GOOG-AT-MDL-B-002764178
145.	GOOG-AT-MDL-019552139	180.	GOOG-AT-MDL-B-002764191
146.	GOOG-AT-MDL-019571201	181.	GOOG-AT-MDL-B-002797051
147.	GOOG-AT-MDL-019588187	182.	GOOG-AT-MDL-B-002803919
148.	GOOG-AT-MDL-019633443		
149.	GOOG-AT-MDL-019642313		
150.	GOOG-AT-MDL-019653406		
151.	GOOG-AT-MDL-019721340		
152.	GOOG-AT-MDL-019767203		
153.	GOOG-AT-MDL-B-000134141		
154.	GOOG-AT-MDL-B-001084151		
155.	GOOG-AT-MDL-B-001140202		
156.	GOOG-AT-MDL-B-002087955		
157.	GOOG-AT-MDL-B-002088697		
158.	GOOG-AT-MDL-B-002088752		
159.	GOOG-AT-MDL-B-002088926		
160.	GOOG-AT-MDL-B-002090567		

HIGHLY CONFIDENTIAL

183.	GOOG-AT-MDL-B- 002837096	205.	GOOG-AT-MDL-B- 006365981
184.	GOOG-AT-MDL-B- 003181628	206.	GOOG-AT-MDL-B- 006939056
185.	GOOG-AT-MDL-B- 003735823	207.	GOOG-AT-MDL-B- 007212533
186.	GOOG-AT-MDL-B- 003741154	208.	GOOG-AT-MDL-B- 007229334
187.	GOOG-AT-MDL-B- 004008753	209.	GOOG-AT-MDL-B- 007232867
188.	GOOG-AT-MDL-B- 004015880	210.	GOOG-AT-MDL-B- 007353902
189.	GOOG-AT-MDL-B- 004016318	211.	GOOG-DOJ-12948968
190.	GOOG-AT-MDL-B- 004243337	212.	GOOG-DOJ-13897780
191.	GOOG-AT-MDL-B- 004247242	213.	GOOG-DOJ-13899823
192.	GOOG-AT-MDL-B- 004265772	214.	GOOG-DOJ-13911836
193.	GOOG-AT-MDL-B- 004425247	215.	GOOG-DOJ-13930748
194.	GOOG-AT-MDL-B- 004680051	216.	GOOG-DOJ-13940086
195.	GOOG-AT-MDL-B- 005083974	217.	GOOG-DOJ-14008698
196.	GOOG-AT-MDL-B- 005167304	218.	GOOG-DOJ-14034714
197.	GOOG-AT-MDL-B- 005168118	219.	GOOG-DOJ-14113270
198.	GOOG-AT-MDL-B- 005170475	220.	GOOG-DOJ-14139857
199.	GOOG-AT-MDL-B- 005282318	221.	GOOG-DOJ-14155066
200.	GOOG-AT-MDL-B- 005372599	222.	GOOG-DOJ-14156827
201.	GOOG-AT-MDL-B- 005457387	223.	GOOG-DOJ-14161619
202.	GOOG-AT-MDL-B- 006069467	224.	GOOG-DOJ-14161943
203.	GOOG-AT-MDL-B- 006316352	225.	GOOG-DOJ-14232497
204.	GOOG-AT-MDL-B- 006365895	226.	GOOG-DOJ-14352302
		227.	GOOG-DOJ-14352774
		228.	GOOG-DOJ-14365517
		229.	GOOG-DOJ-14433486
		230.	GOOG-DOJ-14433633
		231.	GOOG-DOJ-14435110
		232.	GOOG-DOJ-14436029
		233.	GOOG-DOJ-14453674
		234.	GOOG-DOJ-14458088
		235.	GOOG-DOJ-14494204
		236.	GOOG-DOJ-14544715
		237.	GOOG-DOJ-14544743
		238.	GOOG-DOJ-14556699
		239.	GOOG-DOJ-14712739
		240.	GOOG-DOJ-14717683
		241.	GOOG-DOJ-14733660
		242.	GOOG-DOJ-14735212

HIGHLY CONFIDENTIAL

243.	GOOG-DOJ-14735427	287.	GOOG-DOJ-16976651
244.	GOOG-DOJ-14740782	288.	GOOG-DOJ-17101302
245.	GOOG-DOJ-14743636	289.	GOOG-DOJ-17105055
246.	GOOG-DOJ-14870370	290.	GOOG-DOJ-17581278
247.	GOOG-DOJ-15021111	291.	GOOG-DOJ-19090482
248.	GOOG-DOJ-15022600	292.	GOOG-DOJ-19119112
249.	GOOG-DOJ-15022611	293.	GOOG-DOJ-
250.	GOOG-DOJ-15029681		27760500_DUP_1
251.	GOOG-DOJ-15041717	294.	GOOG-DOJ-27762649
252.	GOOG-DOJ-15042589	295.	GOOG-DOJ-27762691
253.	GOOG-DOJ-15064786	296.	GOOG-DOJ-27796586
254.	GOOG-DOJ-15071642	297.	GOOG-DOJ-27803156
255.	GOOG-DOJ-15073261	298.	GOOG-DOJ-27803505
256.	GOOG-DOJ-15076754	299.	GOOG-DOJ-27803533
257.	GOOG-DOJ-15140608	300.	GOOG-DOJ-27804205
258.	GOOG-DOJ-15173140	301.	GOOG-DOJ-27804876
259.	GOOG-DOJ-15186471	302.	GOOG-DOJ-28385887
260.	GOOG-DOJ-15278949	303.	GOOG-DOJ-28420330
261.	GOOG-DOJ-15371972	304.	GOOG-DOJ-29427368
262.	GOOG-DOJ-15417170	305.	GOOG-DOJ-29453844
263.	GOOG-DOJ-15426837	306.	GOOG-DOJ-30066802
264.	GOOG-DOJ-15427800	307.	GOOG-DOJ-32018179
265.	GOOG-DOJ-15428154	308.	GOOG-DOJ-32018452
266.	GOOG-DOJ-15432090	309.	GOOG-DOJ-32022422
267.	GOOG-DOJ-15433127	310.	GOOG-DOJ-32062262
268.	GOOG-DOJ-15434544	311.	GOOG-DOJ-32262980
269.	GOOG-DOJ-15445478	312.	GOOG-DOJ-32265694
270.	GOOG-DOJ-15526215	313.	GOOG-DOJ-AT-00003565
271.	GOOG-DOJ-15590044	314.	GOOG-DOJ-AT-00087809
272.	GOOG-DOJ-15595543	315.	GOOG-DOJ-AT-00252559
273.	GOOG-DOJ-15597407	316.	GOOG-DOJ-AT-00569953
274.	GOOG-DOJ-15598151	317.	GOOG-DOJ-AT-00575435
275.	GOOG-DOJ-15601229	318.	GOOG-DOJ-AT-00588995
276.	GOOG-DOJ-15616526	319.	GOOG-DOJ-AT-00651561
277.	GOOG-DOJ-15621021	320.	GOOG-DOJ-AT-00663174
278.	GOOG-DOJ-15628088	321.	GOOG-DOJ-AT-00750354
279.	GOOG-DOJ-15631978	322.	GOOG-DOJ-AT-01019411
280.	GOOG-DOJ-15637110	323.	GOOG-DOJ-AT-01027937
281.	GOOG-DOJ-15675958	324.	GOOG-DOJ-AT-01514374
282.	GOOG-DOJ-15733030	325.	GOOG-DOJ-AT-01525829
283.	GOOG-DOJ-15766965	326.	GOOG-DOJ-AT-01682499
284.	GOOG-DOJ-15789579	327.	GOOG-DOJ-AT-01811246
285.	GOOG-DOJ-16271009	328.	GOOG-DOJ-AT-01814428
286.	GOOG-DOJ-16732920	329.	GOOG-DOJ-AT-01814688

HIGHLY CONFIDENTIAL

330.	GOOG-DOJ-AT-01815482	374.	GOOG-NE-10660658
331.	GOOG-DOJ-AT-01816686	375.	GOOG-NE-10660882
332.	GOOG-DOJ-AT-01848255	376.	GOOG-NE-10777158
333.	GOOG-DOJ-AT-01933085	377.	GOOG-NE-10952313
334.	GOOG-DOJ-AT-02070545	378.	GOOG-NE-10976657
335.	GOOG-DOJ-AT-02122748	379.	GOOG-NE-10977717
336.	GOOG-DOJ-AT-02146896	380.	GOOG-NE-10978402
337.	GOOG-DOJ-AT-02153612	381.	GOOG-NE-10985565
338.	GOOG-DOJ-AT-02275074	382.	GOOG-NE-11813912
339.	GOOG-DOJ-AT-02275933	383.	GOOG-NE-11847370
340.	GOOG-DOJ-AT-02319393	384.	GOOG-NE-11877282
341.	GOOG-DOJ-AT-02482767	385.	GOOG-NE-12727978
342.	GOOG-DOJ-AT-02634336	386.	GOOG-NE-12738745
343.	GOOG-DOJ-AT-02639830	387.	GOOG-NE-12742008
344.	GOOG-NE-01787563	388.	GOOG-NE-12787759
345.	GOOG-NE-02557667	389.	GOOG-NE-12885078
346.	GOOG-NE-02558055	390.	GOOG-NE-13004102
347.	GOOG-NE-03231727	391.	GOOG-NE-13200100
348.	GOOG-NE-03236272	392.	GOOG-NE-13205235
349.	GOOG-NE-03467508	393.	GOOG-NE-13205865
350.	GOOG-NE-03616222	394.	GOOG-NE-13210272
351.	GOOG-NE-03627597	395.	GOOG-NE-13327192
352.	GOOG-NE-03727939	396.	GOOG-NE-13340752
353.	GOOG-NE-03844060	397.	GOOG-NE-13349343
354.	GOOG-NE-03869994	398.	GOOG-NE-13367768
355.	GOOG-NE-03872605	399.	GOOG-NE-13369624
356.	GOOG-NE-04304005	400.	GOOG-NE-13373170
357.	GOOG-NE-04689402	401.	GOOG-NE-13374150
358.	GOOG-NE-05241137	402.	GOOG-NE-13379438
359.	GOOG-NE-05270679	403.	GOOG-NE-13386334
360.	GOOG-NE-06113062	404.	GOOG-NE-13389481
361.	GOOG-NE-06547825	405.	GOOG-NE-13390045
362.	GOOG-NE-06568562	406.	GOOG-NE-13393115
363.	GOOG-NE-06568719	407.	GOOG-NE-13401911
364.	GOOG-NE-06729717	408.	GOOG-NE-13405025
365.	GOOG-NE-06841582	409.	GOOG-NE-13415537
366.	GOOG-NE-07841238	410.	GOOG-NE-13541543
367.	GOOG-NE-07946731	411.	GOOG-NE-13569193
368.	GOOG-NE-08116078	412.	GOOG-NE-13577462
369.	GOOG-NE-08112779	413.	GOOG-NE-13589899
370.	GOOG-NE-09138394	414.	GOOG-NE-13624783
371.	GOOG-NE-09698556	415.	GOOG-TEX-00039266
372.	GOOG-NE-09711564	416.	GOOG-TEX-00079344
373.	GOOG-NE-10573952	417.	GOOG-TEX-00089835

HIGHLY CONFIDENTIAL

418.	GOOG-TEX-00090151	462.	GOOG-TEX-00806640
419.	GOOG-TEX-00090282	463.	GOOG-TEX-00806682
420.	GOOG-TEX-00090969	464.	GOOG-TEX-00814407
421.	GOOG-TEX-00092657	465.	GOOG-TEX-00825713
422.	GOOG-TEX-00105202	466.	GOOG-TEX-00828547
423.	GOOG-TEX-00105361	467.	GOOG-TEX-00831660
424.	GOOG-TEX-00106255	468.	GOOG-TEX-00850729
425.	GOOG-TEX-00109167	469.	GOOG-TEX-00858576
426.	GOOG-TEX-00110068	470.	GOOG-TEX-00905673
427.	GOOG-TEX-00110540	471.	GOOG-TEX-00959457
428.	GOOG-TEX-00110544	472.	GOOG-TEX-00959461
429.	GOOG-TEX-00111494	473.	GOOG-TEX-00969525
430.	GOOG-TEX-00116069	474.	GOOG-TEX-00978814
431.	GOOG-TEX-00116639	475.	GOOG-TEX-01004466
432.	GOOG-TEX-00119737	476.	GOOG-TEX-01036150
433.	GOOG-TEX-00119845	477.	GOOG-TEX-01142635
434.	GOOG-TEX-00120626	478.	GOOG-TEX-01244428
435.	GOOG-TEX-00120775	479.	GOOG-TEX-01279945
436.	GOOG-TEX-00156142	480.	GOOG-TEX-00974499
437.	GOOG-TEX-00167119	481.	GOOG-TEX-00001418
438.	GOOG-TEX-00177559	482.	GOOG-TEX-00597317
439.	GOOG-TEX-00206687		
440.	GOOG-TEX-00216163		
441.	GOOG-TEX-00234150		
442.	GOOG-TEX-00240572		
443.	GOOG-TEX-00270127		
444.	GOOG-TEX-00271177		
445.	GOOG-TEX-00309326		
446.	GOOG-TEX-00326649		
447.	GOOG-TEX-00336527		
448.	GOOG-TEX-00344083		
449.	GOOG-TEX-00370306		
450.	GOOG-TEX-00374779		
451.	GOOG-TEX-00375239		
452.	GOOG-TEX-00452866		
453.	GOOG-TEX-00513684		
454.	GOOG-TEX-00643890		
455.	GOOG-TEX-00689539		
456.	GOOG-TEX-00705131		
457.	GOOG-TEX-00715805		
458.	GOOG-TEX-00716782		
459.	GOOG-TEX-00777573		
460.	GOOG-TEX-00778301		
461.	GOOG-TEX-00797340		

APPENDIX D

UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF TEXAS
SHERMAN DIVISION

The State of Texas, et. al.

Plaintiff,

v.

Google LLC,

Defendant.

Case No: 4:20-cv-00957-SDJ

Expert Report of Matthew Weinberg

6/7/2024



Matthew Weinberg

and increased publisher revenue, which is expected from an auction theory perspective, as explained above.

100. The amount of data collected is another benefit of a simultaneous auction format like header bidding over a sequential format like waterfalling. Specifically, in a simultaneous auction format, the header bidding setup records the bids of every exchange whereas in a sequential auction format the ad server learns bids of only the winning exchange and those before it in the waterfall. As noted in previous sections, data is crucial in auctions in order to set optimal reserves.^{104, 105}

IV. CONDUCT ANALYSIS: DYNAMIC ALLOCATION AND ENHANCED DYNAMIC ALLOCATION

101. In this section, I analyze the conduct referred to as Dynamic Allocation and later updated to Enhanced Dynamic Allocation.¹⁰⁶ In this section, I draw conclusions regarding their effects. I find that Dynamic Allocation led to a higher win rate¹⁰⁷ and higher revenue for AdX as well as a lower win rate and lower revenue for non-Google exchanges. Furthermore, Enhanced Dynamic Allocation led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers. Reducing the value of direct deals for advertisers should decrease the revenue earned by publishers via direct deals.

102. Moreover, any related conduct that causes AdX to clear publisher-set reserves more often exacerbates many of my conclusions. Dynamic Revenue Sharing is one such conduct, which I describe in Section VII.¹⁰⁸ Throughout, I briefly note conclusions whose magnitude is increased due to Dynamic Revenue Sharing and related conducts.

¹⁰⁴ This is especially true in ad auctions, due to the fact that each impression is unique, and the task of predicting an advertiser's/exchange's bid for a novel impression must be inferred by their past bids on similar but not necessarily identical impressions.

¹⁰⁵ This particular form of additional data is commonly studied in the sub-field of machine learning called "online learning." Learning every exchange's bid is called "expert feedback" and learning only the winning exchange's bid is called "bandit feedback." It is well-understood, including quantitatively, that expert feedback enables faster and more accurate learning than bandit feedback. See, e.g., Aleksandrs Slivkins. "Introduction to Multi-Armed Bandits" (November 2019). <https://arxiv.org/pdf/1904.07272.pdf> (chapters 5 and 6.)

Header bidding produces expert feedback, while waterfalling produces feedback somewhere between bandit and expert, depending on how late the winning exchange is on the waterfall.

¹⁰⁶ There are many descriptions of Dynamic Allocation and Enhanced Dynamic Allocation, according to the documents I have reviewed. The descriptions I provide here reflect my best understanding. It is possible that at some points in time Dynamic Allocation and Enhanced Dynamic Allocation worked slightly different compared to the descriptions I provide here.

¹⁰⁷ I use the term "win rate" to refer to the number of impressions won divided by the number of impressions made available in the open web display ads market.

¹⁰⁸ Briefly, Dynamic Revenue Sharing is a program where AdX sometimes lowered its take rate in order to clear impressions for which it did not solicit a large enough bid to clear the publisher's price floor plus its standard take rate. See Section VII for detailed description.

A. Dynamic Allocation

103. The impact of Dynamic Allocation on the waterfall process depends on the types of line items present in the waterfall. More specifically, whether there are only static demand sources or there are both static and live demand sources¹⁰⁹ affects how Dynamic Allocation works, as well as its impact on the auction procedure and outcomes. Hence, I analyze Dynamic Allocation separately for static demand sources and live demand sources.

1) Dynamic Allocation with static demand sources

104. I first present an overview of Dynamic Allocation during the period when it was first introduced.¹¹⁰ Initially, all line items were static, so Dynamic Allocation addressed a natural shortcoming of the waterfall format. When all line items competing with AdX are static, **Dynamic Allocation with Static Line Items** adjusts the waterfall process in the following manner:

- a. First, Google's ad server DFP processes the high priority line items¹¹¹ that are not affected by Dynamic Allocation (such as direct deals). If any high priority line item succeeds, the impression is sold, and the waterfall terminates without continuing to subsequent steps.
 - b. Every low priority line item, including AdX, has both a price floor and a Value CPM.¹¹² Next, DFP selects the highest Value CPM among all low priority static line

¹⁰⁹ Throughout this report, I use the term “static line item” when referring to line items that do not correspond to outcomes of any auctions. For example, the line items that Google documentation refers to as sponsorship or standard would be static line items. See Google. “Line item types and priorities.” Accessed on May 31, 2024. <https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en>

¹¹⁰ Dynamic Allocation was introduced by DoubleClick, prior to Google's purchase of the company. DoubleClick documentation from that time points to 2007 as the introduction of Dynamic Allocation. See Google. "DoubleClick Advertising Exchange." Accessed on May 31, 2024.

Google documentation claims it is 2008, while agreeing that it predates Google's acquisition of DoubleClick. GOOG-AT-MDL-008991406 at -6. ("2008—Dynamic Allocation [...] pre doubleclick acquisition.")

¹¹¹ Throughout the report, I use “guaranteed” and “high priority” interchangeably when referring to line items that are at priority 1-10 by Google’s standards. Similarly, I use “non-guaranteed” and “low priority” interchangeably when referring to line items that are at priority 12-16. See Google. “Line item types and priorities.” Accessed on May 31, 2024. <https://web.archive.org/web/20240216154938/https://support.google.com/admanager/answer/177279?hl=en> This is in line with the terminology that Google uses. A Google engineer stated that “███████████”

Declaration of Nitish Korula.”

¹¹² Value CPMs are set by the publishers, and they usually correspond to the value of those line items for the publishers. Google provides the following formula to estimate the value CPM: Value CPM = (Total revenue received from ad tags associated with selected line item/Total number of impressions Ad Manager sent to the selected line

items, that satisfy the targeting criteria for this impression and stores this value as a reserve price r . The publisher sets each Value CPM according to Google's documentation, and equal to the expected value derived from each line item. Because all line items are static, the reserve price then captures the expected maximum value that the publisher can earn by selling the impression to a low priority line item.^{113, 114}

- c. Next, DFP calls AdX to run an auction with a reserve price equal to the maximum of r and AdX's price floor. If AdX returns a price higher than this reserve, the impression is sold through AdX, and the waterfall terminates without continuing to subsequent steps.
- d. Finally, if AdX fails to return a bid higher than this reserve, the impression is sold to the low priority demand source with the highest Value CPM (that satisfy the targeting criteria for this impression).¹¹⁵

item)*1000. Google. "Value CPM." Accessed on May 31, 2024.

<https://web.archive.org/web/20221202071803/https:/support.google.com/admanager/answer/177222?hl=en>

A Google engineer stated "When configuring a remnant line item, a publisher must specify a rate (i.e., a price) for the line item. The publisher may also specify a Value CPM; this might be done when the price does not accurately reflect the value to the publisher of serving the line item. For example, if the publisher gave an advertiser a discount on a line item, the publisher could enter the discounted price as the rate and the undiscounted price as the Value CPM. If a publisher does not specify a Value CPM, then Google sets the Value CPM equal to the rate." and "By configuring multiple line items with different targeting criteria, publishers could configure different Value CPMs based on, for example, time of day or the geography of the relevant user, even for the same demand partner. Some publishers set Value CPMs based on their estimates of what CPM a line item would likely generate (taking into account its historical performance) or based on a fixed price the publisher had negotiated with a particular remnant demand partner. Some publishers set Value CPMs higher than their estimates of what CPM a line item would likely generate to increase competitive pressure in the AdX auction or for other reasons. Under Dynamic Allocation, the Value CPM associated with the best eligible non-guaranteed line item could set the reserve price in the AdX auction." GOOG-AT-MDL-008842393 at -96. August 4, 2023. "Declaration of Nitish Korula."

¹¹³ These static line items can pay more or less than the assigned Value CPM.

¹¹⁴ A sophisticated publisher could ignore Google's suggested formulas and set the Value CPMs however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default. Throughout the text, I use the term "sophisticated publisher" with the following context in mind: Publishers can have varying levels of sophistication when attempting to optimize their revenue in the online display advertising ecosystem. On one end, a 'typical' publisher may set parameters according to their ad server's suggested text without developing a detailed understanding of how those parameters are used. At the other end, a 'sophisticated' publisher may fully digest all available documentation and aim to optimize parameters based on their use case, ignoring suggested text. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation. Furthermore, I use the term "sophisticated advertiser" with the following context in mind: Advertisers have varying levels of sophistication when attempting to optimize their outcomes in the online display advertising ecosystem. On one end, a 'typical' advertiser may trust their ad buying tool to optimize on their behalf and input correct information whenever requested (i.e., a 'typical' advertiser would simply input their correct value for an impression when asked). At the other end, a 'sophisticated' advertiser may fully digest all available documentation and aim to optimize inputs to their ad buying tool based on how these inputs are used, ignoring the ad buying tool's recommendations. They may even be able to optimize while accounting for the possibility of conduct that is never disclosed in publicly available documentation.

¹¹⁵ An overwhelming majority of Google's documentation, and their testimony in GOOG-NE-10780865, supports the given definition. GOOG-NE-10780865 at -81. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech

105. The excerpt in Figure 16 from a public Google report from 2010¹¹⁶ reiterates how Dynamic Allocation with Static Line Items work. The document notes that “[Dynamic Allocation] uses this CPM value as the minimum CPM for the auction,”¹¹⁷ as I described above.

Business.” (“If AdX buyers did not bid above the floor price (on a net basis, i.e., after consideration of AdX’s revenue share), the static remnant line item with the highest fixed or estimated price would win the impression.”) However, some internal documentation supports the following conflicting definition:

Of course, any precise mathematical analysis differs between these two formats, although my relevant conclusions do not qualitatively differ. I focus my analysis on the definition which is overwhelmingly supported by Google’s documentation.


¹¹⁶ Google. “Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers.” Accessed on May 31, 2024.

https://web.archive.org/web/20120130063019/http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf (pg. 3 describes mechanics of AdX under Dynamic Allocation.)

¹¹⁷ Google. “Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers.” Accessed on May 31, 2024.

https://web.archive.org/web/20120130063019/http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf

Figure 16: Excerpt from a public Google report explaining how Dynamic Allocation with Static Line Items works¹¹⁸



Dynamic allocation passes to the Ad Exchange the CPM value associated with the ad that the primary ad server has selected and is about to serve. The technology then uses this CPM value as the minimum CPM for the auction. If the Ad Exchange can provide the publisher with a net CPM value higher than they would have gotten from delivering their directly booked, non-guaranteed ad, the Ad Exchange will deliver an ad. If, however, the directly booked ad's CPM value is higher, it ignores any bids coming in from the Ad Exchange. As a result of this ad server integration, publishers essentially have a risk-free way to get the highest yield for every non-guaranteed impression they sell through their direct and indirect sales channels. An additional benefit of Dynamic Allocation is that it ensures there is a deep pool of ads to deliver for any given piece of inventory, reducing the probability that the publisher delivers a house or zero-value ad.


106. The following excerpt in Figure 17 from an internal Google slide deck¹¹⁹ is also consistent with my explanation on how Dynamic Allocation with static line items works.

¹¹⁸ Google. "Profiting from Non-Guaranteed Advertising: The Value of Dynamic Allocation & Auction Pricing for Online Publishers." Accessed on May 31, 2024.

https://web.archive.org/web/20120130063019/http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/doubleclick/pdfs/DC_Ad_Exchange_WP_100713.pdf (pg. 3 describes mechanics of AdX under Dynamic Allocation.)

¹¹⁹ GOOG-NE-08112779. "PBS Basics Training (3) AdX Basics."

Figure 17: An excerpt from an internal Google slice deck explaining Dynamic Allocation with Static Line Items compares the highest valued line item to the AdX clearing price¹²⁰



107. Several internal Google documents provide details on how Dynamic Allocation worked. An internal Google document states that the highest available bid information is passed on to AdX as a “dynamic floor price.”¹²¹ A publicly available Google document states that “Dynamic Allocation is a unique technology that works by passing to the Ad Exchange the CPM value associated with any non-guaranteed ad that DFP is about to serve...the Ad Exchange only serves ads when it can offer a higher price for ad space.”¹²²

¹²⁰ GOOG-NE-08112779 at -94. “PBS Basics Training (3) AdX Basics.”

¹²¹ GOOG-NE-03597611 at -19. December 15, 2011. “Mysteries of Dynamic Allocation.”

¹²² Google. “Maximizing advertising revenues for online publishers.” Accessed on May 31, 2024.

https://web.archive.org/web/20160911040651/https://static.googleusercontent.com/media/www.google.com/en/googleblogs/pdfs/revenue_maximization_090210.pdf


108. To illustrate how Dynamic Allocation with Static Line Items works, imagine an impression arrives for a user over the age of 25 who likes running. DFP notices that this satisfies the coarse targeting criteria for a direct deal with Altra, displays Altra's ad, and the waterfall (even with Dynamic Allocation) ends here. This example is illustrated in Figure 18 below.

Figure 18: An impression is allocated to an Altra direct deal in Dynamic Allocation because it fulfills the targeting criteria



109. Another new impression arrives for a user over the age of 25 who likes expensive handbags. DFP does not have a direct deal that meets this targeting criteria, and so moves on to lower-priority line items and observes that the highest Value CPM option is from a static line item for \$4. DFP then calls AdX with a reserve of \$4. AdX does not find a buyer above \$4, and the waterfall concludes by allocating the impression to the static line item for \$4. This example is illustrated in Figure 19 below.

Figure 19: An impression is allocated to a non-guaranteed line item because the impression fails to clear the targeting criteria of an Altra direct deal, and AdX fails to clear the Dynamic Allocation reserve¹²³



110. Lastly, an impression arrives for a user over the age of 25 who likes expensive handbags, and with further fine-grained cookies noting that the user has recently visited several handbag-shopping websites and lives in a wealthy neighborhood. DFP does not have a direct deal that meets the coarse targeting criteria of users over the age of 25 who like expensive handbags, and so it moves on to lower priority line items and observes that the highest Value CPM option is from a static line item for \$4. DFP then calls AdX with a reserve of \$4. AdX's auction concludes with Gucci winning at a clearing price of \$6. Gucci wins the impression through AdX, paying \$6.¹²⁴ This example is illustrated in Figure 20 below.

¹²³ In the figures, Dynamic Allocation is referred to as "DA."

¹²⁴ This example abstracts away from the role of ad buying tools, ad exchange fees and a publisher-set reserve price for AdX for clarity.

Figure 20: Gucci wins an impression through AdX because it is able to clear the Dynamic Allocation reserve



111. The waterfall format is a suboptimal mechanism for soliciting bids from live sources.^{125, 126} The optimal format would solicit bids from all sources simultaneously and conclude the auction according to Myerson's (1981)¹²⁷ optimal auction.¹²⁸ In the special case where only AdX (or another single exchange) provides live demand and all other demand sources are static, Myerson's optimal auction is equivalent to posting a reserve price that applies to the single live demand source. Here the reserve price would be informed by expected demand from static sources and the optimal auction would sell to the static demand source with the highest Value CPM when the live demand source failed to exceed the reserve price.¹²⁹ Dynamic Allocation

¹²⁵ By live sources, I mean the demand sources that hold live ad auctions, such as exchanges.

¹²⁶ See Section III.B.1 for further discussion on the suboptimality of the waterfall format.

¹²⁷ Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.

¹²⁸ In the independent private values model, Optimal auctions in the (not necessarily independent) private values, and especially interdependent private values models are significantly more complicated. The optimal format in these more complex settings still solicits bids from all sources simultaneously, and sequential formats are suboptimal.

¹²⁹ When restricted to treating the single live demand source as a waterfall line item, several problems might arise. First, if the live demand source is anywhere but first in the waterfall, the impression might be sold before the live source is queried, who might have returned a higher bid. Second, if the live demand source is somehow first in the waterfall but with a low reserve price, it might win the impression despite returning a bid below the best static line item. Together, these identify that the waterfall process with one live demand source can only be optimal if the live

matches this general framework, but by default sets a suboptimal reserve on AdX (the single live demand source). In particular, the maximum Value CPM is certainly not the optimal reserve. Optimal pricing always sets a reserve strictly greater than the opportunity cost (when the opportunity cost is static and known).^{130, 131, 132}

112. Suboptimality notwithstanding, during a period when AdX was the only live demand source, Dynamic Allocation might lead to an increase in revenue for publishers in comparison to not calling live sources at all. When all other demand sources are static, Dynamic Allocation simply gives the publisher a shot at additional revenue (even if that shot is taken sub-optimally by default). Indeed, internal Google documents introducing Dynamic Allocation present the ability to call a live demand source at all as a primary benefit.¹³³ Later on, however, other live demand sources (such as competing exchanges) arose, and as a result, this argument for the benefit of Dynamic Allocation is no longer valid. That is, the primary arguments supporting Dynamic Allocation with only static line items as potentially leading to higher revenue for publishers do not at all apply when line items are live. I now analyze the impact of Dynamic Allocation when other line items are themselves live.

2) Dynamic Allocation with live demand sources

113. In the years after the implementation of Dynamic Allocation, publishers gained the ability to integrate multiple live demand sources, such as exchanges, to DFP.¹³⁴ With the inclusion of live demand sources, Dynamic Allocation changes the waterfall process in the following manner:

source goes first with a reserve exceeding the Value CPM of the best static line item. Dynamic Allocation with Static Line Items allows a sophisticated publisher to satisfy this property, although by default sets a suboptimal (too low) reserve equal to the maximum Value CPM.

¹³⁰ A second source of suboptimality is that static line items are measured in ‘Value CPM’ which are set by the publishers themselves and might be unequal to the CPM (in which AdX is measured), if the ads are of high/low quality, or if discounts were applied. See Google. “Value CPM.” Accessed on May 31, 2024.

<https://web.archive.org/web/20221202071803/https://support.google.com/admanager/answer/177222?hl=en> Whereas AdX competes directly with CPM, so the comparison of CPM from AdX to Value CPM from static line items is not quite right. This can again be mitigated by a sophisticated publisher.

¹³¹ Because there is only one live demand source, the three models (independent private values, private values, and interdependent values) are identical (because there is no uncertainty about the static bids). Therefore, Myerson’s optimal auction is truly the optimal mechanism with one live demand source and several static demand sources.

¹³² As noted, publishers could increase the reserve set on AdX beyond the maximum Value CPM, which would let them recover the optimal format in this special case with only one live demand source, but the default behavior of Dynamic Allocation sets the maximum Value CPM as AdX’s reserve.

¹³³ GOOG-NE-03597611 at -13. December 15, 2011. “Mysteries of Dynamic Allocation.” (“Dynamic Allocation [...] maximizes publishers’ yield [...] By serving AdX [...] whenever they offer more than the competing booked ad networks (real-time competition).”)

¹³⁴ See Interactive Advertising Bureau. “OpenRTB.” Accessed on June 4, 2024.

<https://web.archive.org/web/20240326073202/https://www.iab.com/guidelines/openrtb/> (OpenRTB protocol predates header bidding and gave the publishers the ability to sell impressions through live demand sources.)

- a. First, DFP processes the high priority line items (such as direct deals). If any high priority line items succeed, the impression is sold, and the waterfall terminates without continuing to subsequent steps.
- b. Every lower priority static line item, including AdX, has both a price floor and a Value CPM. Next, DFP computes the highest Value CPM among all low priority line items that satisfy the targeting criteria for this impression, and stores this as a reserve price r .
 - i. If the relevant line item is static, the Value CPM is equal to the (value-adjusted)¹³⁵ CPM earned by the publisher for selecting that line item.¹³⁶
 - ii. If the relevant line item is an exchange, the Value CPM is equal to the (value-adjusted) average historical CPM paid by this exchange to this publisher for past similar impressions.^{137, 138}
 - iii. If the relevant line item is a header bidding bid,¹³⁹ the Value CPM is equal to the header bid.¹⁴⁰ Note that if other exchanges participate primarily via header bidding, this makes the maximum value CPM equal to the clearing price of the header bidding auction.^{141, 142}

¹³⁵ I use the parenthetical (value-adjusted) to note that publishers may care about more than just the revenue earned (for example due to ad quality, a discount offered, etc.).

¹³⁶ A sophisticated publisher could ignore Google's suggested formulas and set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹³⁷ As noted above, this is how publishers typically set Value CPMs, although they are free to deviate from this default behavior. See GOOG-AT-MDL-008842393 at -96. August 4, 2023. "Declaration of Nitish Korula." ("Up until at least December 2021, publishers could set the CPM for their booked static remnant line items (also referred to as "Value CPMs").")

¹³⁸ A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹³⁹ DFP may not know that the relevant line item is a header bidding bid, only that it is a remnant line item. This does not change any conclusions I make.

¹⁴⁰ Notice that the header bidding bid is neither static nor a historical average, it is the winning bid from the header bidding auction. See Section III.B.2 for more details on how header bidding auctions are conducted.

¹⁴¹ Publishers had the ability to increase the clearing price passed on from their header bidding setup to DFP, such as with a multiplier, or an added value. See Asmaâ Bentahar. "Bid Adjustments Simplified: Run Fair Auctions with no Hassle" (May 2, 2021). Accessed on May 31, 2024.

<https://web.archive.org/web/20231202021004/https://www.pubstack.io/topics/bid-adjustments-simplified> (The webpage explains how to implement "bid adjustments" under the section "What are my current Bid Adjustments, and how do I update them?")

¹⁴² A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.


- c. Next, DFP calls AdX with reserve price equal to the maximum of r and AdX's price floor. If AdX succeeds, the impression is sold through AdX, and the waterfall terminates without continuing to subsequent steps.
- d. Finally, if AdX fails to return a clearing price higher than this reserve, the ad server visits low priority line items one at a time in decreasing order of Value CPM. When a static line item is visited, it immediately clears the impression. When a header bidding line item is visited, it immediately clears the impression. When an exchange line item is visited, it runs a live auction for the impression, which may or may not win, so the waterfall might continue after visiting an exchange line item if the impression does not clear.

114. In particular, the standard static waterfall auction, without applying Dynamic Allocation, proceeds by executing steps (a) and (d) above, in which all demand sources, including AdX, are processed in order of their Value CPM (which is static). Dynamic Allocation with Static Line Items is a special case of Dynamic Allocation without exchange or header bidding line items.

115. An internal Google deck¹⁴³ details the explanation laid out above, an excerpt from which can be seen in Figure 21. It notes that the other demand sources compete based on historical CPM, whereas AdX competes with real-time bids.

¹⁴³ GOOG-DOJ-27769247. September 2, 2016. "Header Bidding and FAN."

Figure 21: An excerpt from an internal Google deck that explains how Dynamic Allocation with live demand sources works¹⁴⁴



The way DFP works, AdX gets access to all the non-guaranteed inventory and price it on a real time basis.

Whereas other exchanges or networks, only compete on the basis of historical average price - which means they don't get to see all the inventory and make a real time bid, leading to lower yields for publishers.

116. To illustrate how Dynamic Allocation with live demand sources work, imagine an impression arrives for a user over the age of 25 who likes running. Before DFP executes, header bidding¹⁴⁵ solicits bids from the exchanges integrated into the publisher's header bidding setup. The highest bidder that exceeds their personalized reserve is Nike for \$4 through OpenX, and this is entered as a line item in DFP. DFP then begins the waterfall and notices that this satisfies the targeting criteria for a direct deal with Altra, and decides to display Altra's ad. The waterfall (even with Dynamic Allocation) terminates here. This example is illustrated in Figure 22 below.

¹⁴⁴ GOOG-DOJ-27769247 at -68. September 2, 2016. "Header Bidding and FAN."


¹⁴⁵ That is, if the publisher set up header bidding auction on their website. For this example, I assume the publisher did.

Figure 22: An impression is allocated to an Altra direct deal under Dynamic Allocation because it satisfies the targeting criteria



117. Another impression arrives for a user over the age of 25 who likes expensive handbags. Before DFP executes, the header bidding auction solicits bids from the header bidding exchanges. The highest bidder that exceeds their personalized reserve is Prada for \$4.50 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall. DFP does not have a direct deal that meets this targeting criteria, and so it moves on to lower priority line items and observes that the highest Value CPM option is Prada's header bidding winning bid of \$4.50. DFP then calls AdX with a reserve of \$4.50. AdX does not find a buyer above \$4.50. The waterfall then concludes by selling the impression to Prada for \$4.50. This example is illustrated in Figure 23 below.


Figure 23: An impression is allocated to the header bidding winner because AdX fails to clear the Dynamic Allocation reserve



118. A final impression arrives for a user over the age of 25 who likes expensive handbags, and with further fine-grained cookies noting that the user has recently visited several handbag-shopping websites, lives in a wealthy neighborhood. Before DFP executes, the header bidding auction solicits bids from the exchanges integrated into the publisher's header bidding setup. The highest bidder that exceeds their personalized reserve is Prada for \$4.50 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall. DFP does not have a direct deal that meets the coarse targeting criteria, and so moves on to lower priority line items and observes that the highest Value CPM option is Prada's header bidding winning bid of \$4.50. DFP then calls AdX with a reserve of \$4.50 and the auction concludes with Gucci winning at a clearing price of \$6. Gucci wins the impression through AdX, paying \$6.¹⁴⁶ This example is illustrated in Figure 24 below.

¹⁴⁶ This example abstracts away from the ad buying tools and ad exchange fees for clarity.

Figure 24: Gucci wins an impression through AdX since it clears the Dynamic Allocation reserve



119. If all exchanges, including AdX, participated in header bidding, as opposed to the Dynamic Allocation process, this would create a standard first-price auction with personalized reserves.¹⁴⁷ When one exchange is singled out for default Dynamic Allocation, that exchange is clearly advantaged.¹⁴⁸ In essence, the auction format in this hypothetical is now: (a) all exchanges except for AdX submit a bid in a sealed bid format without learning any others' bids, (b) AdX sees the current highest bid, and then submits its own, (c) the highest bid wins and pays their bid.¹⁴⁹ Hence, Dynamic Allocation allows AdX (and only AdX) to learn others' bids in a first-price auction format, and as a result, Dynamic Allocation creates information asymmetries that favor Google's AdX.¹⁵⁰ This advantage is often referred to as AdX's *Last Look* advantage. I further overview this concept in Section V.C.

¹⁴⁷ It is potentially a bit more complicated than this, only because the bidders in this auction are exchanges rather than direct buyers, but aside from this complication it would just be a first-price auction with personalized reserves.

¹⁴⁸ If a sophisticated publisher uses Dynamic Allocation's flexibility to boost the AdX reserve, Dynamic Allocation still creates an information asymmetry, although the ultimate impact on AdX is indeterminate. I discuss this in more detail when I discuss Last Look in Section V.C.

¹⁴⁹ In reality, it is more complicated than this, since some line items are static and some are "live", but a static bid is still a bid.

¹⁵⁰ See Section II.C for further discussion on how learning others' bids can create advantages in a first-price auction.

B. The Impact of Dynamic Allocation

120. Based on the above articulation of how Dynamic Allocation works,¹⁵¹ I can draw conclusions on its effects. In my opinion, Dynamic Allocation led to higher win rate and higher revenue for AdX as well as lower win rate and lower revenue for non-Google exchanges.¹⁵² Additionally, if AdX typically transacts ads of lower quality than non-Google exchanges, Dynamic Allocation also led to an increase in the display of lower quality ads.¹⁵³ This is my opinion in aggregate, accounting for the possibility that some publishers chose to use default options while others chose to cleverly set Value CPMs to optimize revenue, and accounting for periods both when other exchanges participated via the waterfall and when other exchanges participated via header bidding. In the subsections below, I draw more precise conclusions for each of these settings.

1) Impact of Dynamic Allocation when other exchanges participate in the waterfall

121. I first compare AdX to exchanges that participate in the waterfall. In this case, Dynamic Allocation has the following effect as compared to no Dynamic Allocation: (a) AdX is always visited first in the waterfall, and (b) AdX's reserve may be increased. I now draw conclusions of these effects.

122. In comparison to exchanges that participate in the waterfall, under Dynamic Allocation, and no matter how a publisher sets Value CPMs, AdX is the only exchange that always has the opportunity to submit a bid on any inventory that is not sold through high priority line items.¹⁵⁴ More specifically, if AdX returns a bid that exceeds its reserve price set by Dynamic Allocation, then no other exchange has the opportunity to submit a bid because the waterfall stops, so therefore, AdX is the only exchange that always has this opportunity.¹⁵⁵

¹⁵¹ For the entirety of this section, when I say 'Dynamic Allocation,' I refer to 'Dynamic Allocation with Live Demand Sources' and not 'Dynamic Allocation with Static Line Items' (which I introduce both for historical context and to develop intuition as a special case).

¹⁵² In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁵³ In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁵⁴ This is only true when other exchanges are in the waterfall. If other exchanges can access header bidding, they all have an opportunity to bid on each impression.

¹⁵⁵ I again note that AdX would be more likely to return a bid that exceeds its Dynamic Allocation reserve price due to conduct such as Dynamic Revenue Sharing that increases the likelihood that AdX clears its publisher-set reserve price.

a) Dynamic Allocation enables AdX to transact more high-value impressions

123. Let me next consider “high-value” impressions, which have relatively high value given the fine-grained targeting data available to live bidders, but cannot be recognized as such merely on the basis of coarse targeting data used to set static reserves.¹⁵⁶ When other exchanges primarily participate via the waterfall, Dynamic Allocation, no matter how a publisher sets Value CPMs, would lead to AdX winning an even greater volume of high-value impressions, and increased revenues from these impressions under Dynamic Allocation compared to no Dynamic Allocation.¹⁵⁷ This is because AdX submits a live bid based on fine-grained targeting data, whereas its static reserve price is set based only on coarse targeting data (no matter how publishers set Value CPMs, they are still static). If AdX is winning a higher volume of these high-value impressions, then other demand sources necessarily win a lower volume of these high-value impressions. This is one instance where the benefit of going first clearly outweighs the cost of a potentially higher reserve.

124. The impact on “typical-value” impressions is less clear-cut. On one hand, under Dynamic Allocation, AdX has the ability to make *all* impressions available first, meaning that AdX will always get an opportunity to solicit bids. On the other, Dynamic Allocation results in a higher reserve on AdX than without Dynamic Allocation, and so AdX’s solicited bids are less likely to clear its reserve. For impressions of higher-than-average value, the prior reasoning still holds; the cost of facing a higher reserve is still minimal (because it is static), whereas the benefit of a first bite is significant (because live bids are likely to be higher than any of the static reserves).¹⁵⁸ For impressions of lower-than-average value, the cost of facing a higher reserve could outweigh the benefit of going first.

2) Impact of Dynamic Allocation when other exchanges participate in header bidding

125. I now compare AdX to exchanges that participate via header bidding. In this case, via Dynamic Allocation, AdX learns information about bids relayed by other exchanges.

¹⁵⁶ To have an example in mind, imagine a male over the age of 25 who likes running, *and has visited ten different shoe company's websites in the last hour*. Static reserves based on coarse targeting information would treat this impression like any other male over the age of 25 who likes running, while live bids will know that this is an unusually high-yield impression.

¹⁵⁷ In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁵⁸ In my opinion, conduct such as Dynamic Revenue Sharing (which causes AdX to clear its publisher-set price floor more often) would lower the bar for an impression to be considered “higher-than-average” value.

- a) Dynamic Allocation enabled AdX to learn the header bidding clearing price

126. In comparison to exchanges that participate via header bidding, under Dynamic Allocation, AdX is the only exchange that learns information about others' bids and passes it on to its bidders. In particular, all exchanges that participate in header bidding must relay bids that are made without any information regarding other exchanges' bids. If the publisher uses default options, AdX can instead relay bids while knowing the maximum bid returned by all header bidding exchanges. I previously noted that this is referred to as AdX's Last Look advantage and constitutes a significant advantage in an auction (see Section V.C).¹⁵⁹

- b) Dynamic Allocation enabled AdX to win impressions by bidding one cent above the header bidding clearing price

127. To illustrate why the information on other participant's bids is useful in a first-price auction, imagine that Bidder One has a value of \$8 and Bidder Two has a value of \$10, and that they participate in a sealed bid first-price auction. Bidder One would surely not submit a bid of \$8, as this guarantees them no gain (either they lose and gain nothing, or they win and pay \$8, which is again leads to no gain). Bidder One does not know what bid Bidder Two will submit, and so does not know how much to shade their bid. Still, Bidder One may form a belief about Bidder Two's behavior, and compute that their optimal bid, given the information they have, is \$4. Bidder Two similarly will not submit a bid of \$10, similarly does not know how much to shade their bid, and perhaps computes that their optimal bid in response to whatever belief they form about Bidder One's behavior is to shade their bid to \$5. In this case, Bidder Two wins, pays \$5, and enjoys a utility of $\$10 - \$5 = \$5$. Instead, imagine that Bidder One learns Bidder Two's bid before submitting their own. Bidder Two still will surely not submit a bid of \$10. In particular, while they do not know Bidder One's value, Bidder Two guesses that Bidder One will likely outbid them by a penny if Bidder One's value exceeds Bidder Two's bid and lose otherwise. Uncertainty over Bidder One's value may still result in Bidder Two optimally shading their bid to \$5. In this case, Bidder One would then submit a bid of \$5.01 to win the auction, paying \$5.01 and enjoying a utility of $\$8 - \$5.01 = \$2.99$. This example demonstrates the mechanics by which a bidder would win more often and gain by seeing others' bids in a first-price auction. I also provide a more detailed example in Appendix D to more precisely convey this principle.

¹⁵⁹ If a sophisticated publisher instead cleverly sets Value CPMs as a function of header bids, then AdX might still infer information about the maximum header bid, and in particular certainly knows that the maximum header bid lies below its reserve.

128. In the interest of tractable intuition, the above example is simplified (there is no reserve, bidders bid directly without an exchange, etc.). The same principles extend to Dynamic Allocation where the publisher uses default options.

- a. The fact that there is no reserve simplifies the math but is not material to the conclusions. As long as the publisher uses default options in Dynamic Allocation, Bidder One in this example can still win by bidding a penny above Bidder Two when they both exceed the reserve.
- b. The fact that the bidders directly bid in the auction instead of going through exchanges is immaterial if those exchanges themselves run first-price auctions.¹⁶⁰,¹⁶¹,¹⁶² Indeed, bidders who bid in Exchange One still learn the maximum bid from Exchange Two, which still better informs their optimal bid-shading.
- c. If bidders instead bid directly through exchanges that run second-price auctions, the benefit is even starker. A second-price auction executed by Exchange One is truthful, even given the existence of Exchange Two. More importantly, given that Exchange Two's second highest bid sets the reserve for Exchange One, Exchange One will win whenever their highest value exceeds Exchange Two's second highest bid.¹⁶³ On the other hand, Exchange Two wins only when its second highest bid exceeds the highest bid of Exchange One.

¹⁶⁰ Many exchanges, such as AppNexus, Index Exchange and OpenX gradually switched to the first-price auction format. See Sarah Sluis. "Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price" (September 5, 2017). Accessed on May 31, 2024.

<https://web.archive.org/web/20220712083559/https://www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/>

¹⁶¹ A Google engineer stated that "Many exchanges began to move from second-price to first-price auctions in the mid-to-late-2010s." GOOG-AT-MDL-008842393 at -95. August 4, 2023. "Declaration of Nitish Korula."

¹⁶² AdX later switched to the first-price auction format as well, in 2019. See Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024.

<https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/>

¹⁶³ When an exchange uses a second-price auction, this auction itself is truthful. However, without a Last Look advantage, the clearing price of this exchange is then entered into a subsequent first-price auction via header bidding. Therefore, the advertiser has to determine not only their own bid, but ways in which they can affect the clearing price (for example, they might want their clearing price to be higher so that their bid wins in the header bidding auction). While their own winning bid cannot affect the clearing price, a very sophisticated advertiser could create another account and "second-price themselves" in order to increase the clearing price they pay and improve their chances of winning the subsequent first-price auction. All of this is to say that bidding in a truthful second-price auction when that auction is run by an exchange that needs to submit the winning bid to a subsequent first-price auction is complicated. On the other hand, bidding in a second-price auction in an exchange with Last Look advantage is still truthful. As long as a bidder submits a bid that exceeds the maximum bids returned by all other exchanges, and all other advertisers who participate in AdX, that bidder will win (and the price will be their minimum bid to win). Therefore, an exchange running a second-price auction with Last Look advantage provides an advantage to its bidders as compared to an exchange running a second-price auction without a Last Look advantage.

129. I therefore conclude that, when exchanges participate primarily via header bidding, AdX would win more impressions under Dynamic Allocation with default options as compared to no Dynamic Allocation.¹⁶⁴ Because AdX earns revenue via take-rates on cleared transactions, this would cause AdX's revenue to increase. Because this conclusion applies to any impression, it applies to high-value impressions as well. By the same logic, non-Google exchanges would necessarily clear fewer transactions (including high-value transactions) and earn less revenues.¹⁶⁵

3) Impact of Dynamic Allocation on ad quality

130. If it is the case that AdX typically transacts ads that are of lower quality compared to non-Google Exchanges, then an increased win rate for AdX would result in lower quality ads for the publisher.¹⁶⁶ Based on my above conclusions, Dynamic Allocation, no matter how publishers set Value CPMs, would result in lower quality ads displayed on high-value impressions when exchanges participate primarily through the waterfall. Similarly, Dynamic Allocation with default options would result in lower quality ads displayed on all impressions when exchanges participate primarily through header bidding. This follows because Dynamic Allocation results in AdX winning additional impressions, and displaying ads of whatever quality it tends to display.

C. Enhanced Dynamic Allocation

131. Google later amended Dynamic Allocation to include guaranteed line items as well. This is referred to as **Enhanced Dynamic Allocation** and it changes the waterfall process in the following manner:

- a. First, DFP calculates a "temporary CPM"¹⁶⁷ for all high priority line items according to the campaign goals associated with those line items. For example, if the delivery is behind for a specific high priority line item, the temporary CPM associated with

¹⁶⁴ In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁶⁵ The impact of Dynamic Allocation with sophisticated publishers who cleverly set Value CPMs is less clear-cut. On one hand, if sophisticated publishers only slightly inflate the Value CPM of the winning header bid, then the above conclusions continue to hold for exactly the same reasons. On the other hand, if sophisticated publishers significantly inflate the Value CPM of the winning header bid due to Dynamic Allocation and would not have set such an inflated reserve on AdX in absence of Dynamic Allocation, then the cost of this inflated reserve might outweigh the benefits highlighted above.

¹⁶⁶ In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁶⁷ GOOG-AT-MDL-008842393 at -99. August 4, 2023. "Declaration of Nitish Korula." ("Up to at least December 2021, with Enhanced Dynamic Allocation, the ad server calculated what is known as a temporary CPM for a guaranteed deal.")

that line item goes up.¹⁶⁸ DFP stores the highest temporary CPM as a reserve price R .

- b. Next, DFP computes the highest Value CPM among all low priority line items that satisfy the targeting criteria for this impression, and stores this as a reserve price r .
 - i. If the relevant line item is static, the Value CPM is equal to the (value-adjusted) CPM earned by the publisher for selecting that line item.¹⁶⁹
 - ii. If the relevant line item is an exchange, the Value CPM is equal to the (value-adjusted) average historical CPM paid by this exchange to this publisher for past similar impressions.¹⁷⁰
 - iii. If the relevant line item is a header bidding bid, the Value CPM is equal to the header bid. Note that if other exchanges participate primarily via header bidding, this makes the maximum value CPM equal to the clearing price of the header bidding auction.¹⁷¹
- c. Next, DFP calls AdX with reserve price equal to the maximum of r , R , and AdX's price floor. If AdX succeeds, the impression is sold through AdX, and the waterfall terminates without continuing to subsequent steps.
- d. Finally, if AdX fails to return a clearing price higher than this reserve,¹⁷²

¹⁶⁸ A Google engineer explained in a declaration that [REDACTED]

[REDACTED] GOOG-AT-MDL-008842393 at -99. August 4, 2023.

"Declaration of Nitish Korula."

¹⁶⁹ A sophisticated publisher could ignore Google's suggested formulas and set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹⁷⁰ A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹⁷¹ A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹⁷² GOOG-AT-MDL-008842393 at -00. August 4, 2023. "Declaration of Nitish Korula." ("if the highest effective AdX bid could beat both the EDA price and the price of the remnant line item that was selected as a candidate for the impression, then the ad associated with that AdX bid would win. If not, the guaranteed or remnant line item would win.") This suggests that the winner in this case is determined by a comparison between r and R and does not

- i. If $R > r$, the impression is sold to the high priority line item with the highest temporary CPM (equal to R).
- ii. Otherwise, the impression is offered to the low priority line item with the highest Value CPM (equal to r). If this line item is static or a header bid, the impression is immediately cleared. If this line item is an exchange, the exchange runs a live auction with its assigned price floor, and the impression may or may not clear.


132. The key difference between Dynamic Allocation and Enhanced Dynamic Allocation is that high priority line items no longer retain their absolute priority over AdX, which is demonstrated in step c above. Step d describes what happens if AdX fails to clear its reserve.

133. The following excerpt in Figure 25 from an internal Google document on Google auction adjustments¹⁷³ outlines how the reserve price is determined by Enhanced Dynamic Allocation. It corroborates the explanation I presented at the beginning of this subsection.

explicitly state what happens in case the remnant line item is an exchange that does not clear the impression. The Google documentation I have access to does not explicitly clarify this aspect, but exactly how this case resolves is immaterial to my conclusions.

¹⁷³ GOOG-NE-13203009. "DRX Global Optimization of DRS, RPO, and EDA."

Figure 25: An excerpt from an internal Google slide deck on Enhanced Dynamic Allocation, outlining that the dynamic reserve price is a calculation based on both the delivery of the guaranteed line items as well as the maximum valued non-guaranteed line item¹⁷⁴



134. To illustrate how Enhanced Dynamic Allocation impacts the auction process, imagine an impression arrives for a user over the age of 25 who likes running. Before DFP executes, the header bidding auction solicits bids from exchanges they integrated to their header bidding system. The highest bidder that exceeds their personalized reserve is Nike for \$4 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a Value CPM of \$10, which is the highest Value CPM among all line items. DFP then calls AdX with a reserve price of \$10,¹⁷⁵ but AdX does not find a buyer above \$10. The waterfall then concludes by using the impression to fulfill Altra's direct deal. This example is illustrated in Figure 26 below.

¹⁷⁴ GOOG-NE-13203009 at -16. "DRX Global Optimization of DRS, RPO, and EDA."

¹⁷⁵ This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the "temporary CPM" adjustment for the sake of clarity.

Figure 26: An impression is allocated to an Altra direct deal because AdX fails to clear the Enhanced Dynamic Allocation reserve¹⁷⁶



135. Another impression arrives for a user over the age of 25 who likes running, and with further fine-grained cookies noting that the user has recently visited several shoe-shopping websites, lives in a wealthy neighborhood, and competes in races. Before DFP executes, the header bidding auction solicits bids from the exchanges they integrated to their header bidding system. The highest bidder that exceeds their personalized reserve is Nike for \$8 through OpenX, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a Value CPM of \$10, which is the highest Value CPM among all line items. DFP then calls AdX with a reserve of \$10,¹⁷⁷ whose auction concludes with Reebok winning at a clearing price of \$15. Reebok wins the impression through AdX, paying \$15.¹⁷⁸ This example is illustrated in Figure 27 below.

¹⁷⁶ This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the “temporary CPM” adjustment for the sake of clarity.

¹⁷⁷ This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the “temporary CPM” adjustment for the sake of clarity.

¹⁷⁸ This example abstracts away from the ad buying tools and ad exchange fees for clarity.

Figure 27: Reebok wins an impression through AdX because it is able to clear the Enhance Dynamic Allocation reserve set by the Altra direct deal¹⁷⁹



136. These examples show that Enhanced Dynamic Allocation gives AdX the opportunity to bid against high priority line items such as direct deals, too. The last example also highlights one reason that AdX might possibly outbid high priority line items which typically have a much higher CPM compared to remnant line items,¹⁸⁰ that the fine-grained targeting criteria may alert live demand sources of a particularly high value impression and bid higher than the direct deal price that is typically based on coarser targeting criteria.

D. The Impact of Enhanced Dynamic Allocation

137. With the above articulation of how Enhanced Dynamic Allocation works, I now draw conclusions on its effects. In my opinion, Enhanced Dynamic Allocation likely led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers, which would in turn decrease the revenue earned by publishers via direct deals.¹⁸¹ I explain all these conclusions below.

¹⁷⁹ This example assumes that the guaranteed line item price directly becomes the Enhanced Dynamic Allocation without the “temporary CPM” adjustment for the sake of clarity.

¹⁸⁰ Google states that “On average, direct ads sell for two to four times higher than programmatic ads. Direct ads average \$10-20 CPMs (cost per thousand impressions), while programmatic ads average \$1-5 CPMs.” Google. “Understand Direct and Programmatic Ad Revenue.” Accessed on May 31, 2024.

<https://web.archive.org/web/20231226200704/https://newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/>

¹⁸¹ In my opinion, the magnitude of AdX’s win rate and revenue increase would be larger due to conducts such as Dynamic Revenue Sharing which increase the likelihood that AdX clears its publisher-set reserve. The impact of such conducts on direct deals is less clear-cut, and so I do not opine on the further impact of such conducts on direct deals beyond the immediate impact of Enhanced Dynamic Allocation.

138. Enhanced Dynamic Allocation expands Dynamic Allocation to also include high-priority line items. In particular, the ‘Dynamic Allocation Framework’ can be interpreted as (a) finding a set of line items, then (b) allowing AdX to be considered before any of those line items,¹⁸² but (c) setting the reserve price of AdX to be equal to the maximum Value/Temporary CPM of those line items.¹⁸³

139. Under Enhanced Dynamic Allocation, impressions that otherwise would have been reserved for high priority line items such as direct deals are instead available for AdX’s auction. This follows immediately from the definition of Enhanced Dynamic Allocation since it allows AdX to run an auction for all impressions. The Enhanced Dynamic Allocation-generated reserve price might be high, but AdX will always have the opportunity to run an auction for all impressions. Without Enhanced Dynamic Allocation, any impression with a viable high priority line item would not be available to AdX for auction.

140. Furthermore, AdX is the only exchange that unconditionally has this opportunity. Under Enhanced Dynamic Allocation, another exchange *can* have this opportunity, but only if (a) its Value CPM exceeds the highest temporary CPM among high priority line items, and (b) AdX fails to clear its reserve. In particular, (a) suggests a high barrier to this exchange being considered in front of the high priority line item at all,¹⁸⁴ and (b) notes that AdX still gets a first bite, even if the Value CPM of an exchange is high enough to satisfy (a).

141. As Enhanced Dynamic Allocation runs live auctions for every impression, it will likely create a revenue increase for the publishers in the short run.^{185, 186} This conclusion is supported by an internal Google slide deck,¹⁸⁷ an excerpt from which is reproduced in Figure 28. It shows that Enhanced Dynamic Allocation led to an increase in publisher revenue from AdX in the first

¹⁸² When one of these line items is a header bidding line item, Dynamic Allocation considers AdX before that line item, although of course that line item itself was generated from a live bid in an auction executed before the ad server.

¹⁸³ Under default publisher behavior. As previously noted, a sophisticated publisher could further increase AdX’s reserve beyond the maximum Value CPM.

¹⁸⁴ Recall that “On average, direct ads sell for two to four times higher than programmatic ads. Direct ads average \$10-20 CPMs (cost per thousand impressions), while programmatic ads average \$1-5 CPMs.” Google. “Understand Direct and Programmatic Ad Revenue.” Accessed on May 31, 2024.

<https://web.archive.org/web/20231226200704/https://newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/>

This suggests that temporary CPMs of high priority line items would likely be higher than static Value CPMs of low priority line items.


¹⁸⁵ In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁸⁶ Importantly, it is possible that many publishers deprioritize short term revenue, and care more about their revenue in the long run. I elaborate on the long run effects, namely the “cream-skimming effect,” below.

¹⁸⁷ GOOG-NE-03872763. “Discussion on improving AdX & AdSense backfill.”

quarter of 2014. This also makes sense, as high priority line items are all static. I previously noted in Section IV.A that it would increase revenue to allow a live demand source the option to outbid a static demand source.

Figure 28: An excerpt from an internal Google slide deck plotting the impact of Enhanced Dynamic Allocation on publisher metrics¹⁸⁸



142. Since AdX generates revenue by taking a cut of the clearing prices, the plot above shows that Enhanced Dynamic Allocation increases AdX and Google revenue as well. This can be seen by the upward trends of the plots in the figure above. If impressions that satisfy targeting criteria for direct deals are on average more valuable than impressions that do not,¹⁸⁹ then Enhanced

¹⁸⁸ GOOG-NE-03872763 at -85. "Discussion on improving AdX & AdSense backfill."

¹⁸⁹ Google's online documentation states that "Direct ads average \$10-20 CPMs (cost per thousand impressions), while programmatic ads average \$1-5 CPMs." See Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024.

<https://web.archive.org/web/20231226200704/https://newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/>

Dynamic Allocation results in more valuable transactions being transacted through AdX rather than direct deals. Therefore, Enhanced Dynamic Allocation would not only lead to increased volume and revenues for AdX, but also to a greater volume of valuable impressions being transacted through AdX.¹⁹⁰

143. This outcome reduces the value of direct deals for advertisers, which would likely decrease the revenue that publishers can expect to earn via direct deals. Specifically, assuming that the value for impressions is refined based on fine-grained targeting data used in live ad auctions after the coarse targeting criteria used in direct deals, live demand sources are better informed on the value of an impression than direct deal partners for that impression. For example, Altra's value for an impression to a man over the age of 25 who likes running might be \$10 on average, but sometimes \$15 if that user has also visited running shoe websites in the last two hours, and sometimes \$5 if that user has not. In the absence of Enhanced Dynamic Allocation and any cream-skimming effect,¹⁹¹ Altra would enjoy a CPM value of \$10 on average for men over the age of 25 who like running, half of whom have also visited running shoe websites in the last two hours and half of whom who have not. But importantly, Altra's direct deal is based on the coarse targeting criteria of men over the age of 25 who like running. With Enhanced Dynamic Allocation, AdX solicits bids based on the full fine-grained targeting criteria. Therefore, if the Temporary CPM of Altra's direct deal is \$10, AdX is likely to find a bid exceeding its reserve price for users who have recently visited running shoe websites, and not for those who have not. Therefore, Altra's direct deal would exclusively serve low value impressions to the advertisers, coming from users who have not also visited running shoe websites recently and give a value of \$5. Note that this does not limit the publisher's ability to fulfill Altra's direct deal. It simply causes Altra's direct deal to be filled exclusively with low-value impressions rather than a mix of low- and high-value impressions.

144. This cream-skimming effect reduces the value of direct deals with publishers in the long run from the perspective of the advertisers. Hence, it potentially leads to a negative impact on the

¹⁹⁰ In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.

¹⁹¹ In this example, the "cream" of the impressions to users over the age of 25 who like running are those who have visited running shoe websites in the last two hours. AdX "skims the cream" of these users because it is aware of which users are the cream and which are not when soliciting live bids. Altra, on the other hand, is stuck only with the leftovers (despite the fact that Altra submitted a bid generically on users over the age of 25 who like running, half of whom are the "cream"). This is also an instance of what is called *adverse selection*, because Altra might initially think that they are getting a representative impression among users over the age of 25 (for which their average value is \$10), but a better-informed party causes them to receive non-representative impressions for users over the age of 25 (and in particular, the low value ones).

publisher revenue in the long run. There are some means to partially mitigate this impact, but in any auction where Enhanced Dynamic Allocation causes a different outcome than Dynamic Allocation, there is risk of cream-skimming. The key property to consider is that the following two impressions have different values: (a) an impression for an average user who satisfies the coarse targeting criteria, and (b) an impression for an average user who satisfies the coarse targeting criteria and for which AdX does not find a bid exceeding its reserve of r . If r is small, then the values of (a) and (b) might be very far apart. If r is huge, then the values of (a) and (b) might be essentially (or even exactly) the same. Therefore, one mitigation might be to always set temporary CPMs so large that AdX would never find a bid. This is in fact a full mitigation of this effect but results in the same outcome as Dynamic Allocation (and therefore fully nullifies Enhanced Dynamic Allocation). One form of partial mitigation might be to have a reserve price sometimes equal to the true Value CPM of the direct deal, and sometimes so huge that AdX would never find a bid. In the former case, Enhanced Dynamic Allocation results in distinct outcomes from Dynamic Allocation, but the direct deal suffers the full effects of cream-skimming. In the latter, the direct deal does not suffer from cream-skimming, but Enhanced Dynamic Allocation is also just equal to Dynamic Allocation. Another form of partial mitigation might be to have a reserve price higher than the true Value CPM of the direct deal, but not so huge that AdX would never find a bid. In this case, the direct deal suffers some cream-skimming (they only miss out on the exceptionally-high-value impressions), and Enhanced Dynamic Allocation yields only some distinct outcomes from Dynamic Allocation (AdX only has a serious shot at the exceptionally-high-value impressions). These are examples of possible mitigations. For any cream-skimming-mitigation approach, there is a direct tradeoff between how distinct the outcomes are under Enhanced Dynamic Allocation and how much the Direct Deal would be hurt by cream-skimming.

V. CONDUCT ANALYSIS: EXCHANGE BIDDING

145. In this section, I analyze header bidding within the context of auction design, focusing on auction structure and dynamics both prior and subsequent to Google's introduction of Exchange Bidding. My analysis compares header bidding to a waterfall approach (see Section III.B for details) and header bidding to Exchange Bidding, which I discuss in detail below. I find that header bidding improves publisher outcomes relative to the waterfall approach (with and without Dynamic Allocation and Enhanced Dynamic Allocation). I further find that header bidding can generate higher revenue for publishers than can Exchange Bidding.

A. Exchange Bidding

146. I first outline Google's rival product to header bidding. In 2018, Google released **Exchange Bidding**, a technology similar to header bidding, after observing the rise of header bidding technology.^{192, 193, 194} Exchange Bidding created a separate auction of auctions into which all exchanges except for AdX submit their bids. AdX would then submit a bid against the prevailing bid from this auction of auctions. Similar to header bidding, Exchange Bidding was a first-price auction.¹⁹⁵ Furthermore, publishers could use both header bidding and Exchange Bidding at the same time.

147. Exchange Bidding augments the waterfall process in the following manner:¹⁹⁶

a. [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

b. [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

¹⁹² First the tool was called "Exchange Bidding Dynamic Allocation," later renamed to "Exchange Bidding," then to "Open Bidding" after the co-rollouts of Unified Pricing Rules and the AdX auction format change to the first-price. See AdExchanger. "Google's Exchange Bidding Is Now 'Open Bidding'; Market Researchers Slip" (August 27, 2019). Accessed on May 31, 2024. <https://web.archive.org/web/20220523024855/https://www.adexchanger.com/ad-exchange-news/tuesday-27082019/>

The internal development code name for the product was "Jedi." See GOOG-NE-03995243 at -3. July 25, 2018. "PRD: Unified 1P auction and Pricing rules."

¹⁹³ To be more precise, the best product to compare to Exchange Bidding is the header bidding variant called "server-side header bidding," since both handle the auction on a server rather than the user's internet browser. See Anthony Vargas. "AdExplainer: Client-Side vs. Server-Side Header Bidding: What's The Difference?" (December 1, 2023). Accessed on May 31, 2024.

<https://web.archive.org/web/20240314163210/https://www.adexchanger.com/adexplainer/adexplainer-client-side-vs-server-side-header-bidding-whats-the-difference/> ("With client-side header bidding, the bulk of that processing occurs on the user's device in the web browser itself. With server-side header bidding, the processing happens on a remote server.")

¹⁹⁴ Google's First Am. Resp. and Obj. to Plaintiff's Third Set of Interrogs. (May 24, 2024) at 11.

¹⁹⁵ Internal Google documents state that "EB was our first attempt at running a 1P [first-price] auction; Since other exchanges already have experience with submitting 1P bids into HB wrappers, it was the easiest way to build out the product." GOOG-NE-13494966 at -71. May 2019. "Managing Yield."

¹⁹⁶ GOOG-TEX-00000744. April 26, 2017. "Exchange Bidding (Jedi) Open Beta Sates Readiness Review." (internal Google slide deck that discusses how Exchange Bidding worked during its initial implementation, as well as Google's plans on Exchange Bidding rollout.)

i. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

¹⁹⁷ A sophisticated publisher could ignore Google's suggested formulas and set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹⁹⁸ A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default.

¹⁹⁹ A sophisticated publisher could set the Value CPM however they like. If the publisher is sophisticated and revenue-maximizing, they would choose a Value CPM above the default. Recall that some publishers were indeed sophisticated and chose to do this. This is referred to as the "boost." See, e.g., GOOG-TEX-00843142 at -46. September 3, 2019. "First-price bidding Update." ("Why does "boost" exist? The publisher inflates the HB bid before sending it as a floor to AdX. Drives up cost + "fairer" comparison between Google buyer bid and HB bid.")

²⁰⁰ When integrating an exchange into Exchange Bidding, the publisher has the option to set a personalized price floor.

²⁰¹ [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

148. Exchange Bidding can be interpreted as affording some of the advantages provided to AdX to other exchanges,²⁰⁵ since it extends the Enhanced Dynamic Allocation Last Look advantage to them, provided they pay the Exchange Bidding fee. In fact, the earliest iteration of Exchange Bidding was called “Exchange Bidding in Dynamic Allocation.”²⁰⁶

149. To illustrate the process of Exchange Bidding, suppose that there is opportunity to serve an impression for a user over the age of 25 who likes running, and with further fine-grained data from cookies noting that the user has recently visited several shoe-shopping websites, lives in a wealthy neighborhood, and competes in races. Before DFP executes, header bidding solicits bids from exchanges. The highest bidder that exceeds their personalized reserve is Nike for \$8, and this is entered as a line item in DFP, which then begins the waterfall, and notices that this satisfies the targeting criteria for a direct deal with Altra, with a temporary CPM of \$10, which is the highest temporary or Value CPM among all line items. DFP then calls Exchange Bidding and AdX, with a reserve of \$20 for AdX, in order to adjust for the fact that the publisher believes AdX ads to be of low quality, and a reserve of \$10 for Index Exchange.²⁰⁷ AdX’s auction concludes with no bidder exceeding the reserve, and Index Exchange concludes with Vibram²⁰⁸

²⁰² GOOG-TEX-00105361 at -67. April 28, 2017. “FAN Bidding in to DPI and AdMob.”

²⁰³ The bids submitted to the Exchange Bidding auctions are net of this fee, so the amount paid by the advertiser is equal to the bid submitted by the ad buying tool into the exchange plus the ad buying tool fee.

²⁰⁴ Later, along with the introduction of Unified Pricing Rules, Google switched to what was called “Unified Auction” where all exchanges and all “authorized buyers” participated in a single first-price auction with unified reserves. Since the unified auction was a first-price auction with non-personalized reserves, this meant that AdX effectively switched to first-price auction as well. GOOG-NE-13494966 at -80. May 2019. “Managing Yield.”

Under Unified Pricing Rules, publishers cannot set personalized reserve prices on each exchange and must instead set the same reserve price for all exchanges in Exchange Bidding as well as AdX (but this reserve may still exceed r). I analyze Unified Pricing Rules in Section VI.

²⁰⁵ GOOG-DOJ-AT-01809483 at -84. March 2017. “Exchange Bidding in Dynamic Allocation (fka Project Jedi).” (“Exchange Bidding in Dynamic Allocation is a project to allow publishers to invite existing third party exchange partners to bid in real-time alongside AdX in Dynamic Allocation.”)

²⁰⁶ See Jonathan Bellack. “Improving yield, speed and control with DoubleClick for Publishers First Look and exchange bidding” (April 13, 2016). Accessed on May 31, 2024.


<https://web.archive.org/web/20240206070512/https://blog.google/products/admanager/improving-yield-speed-and-control-with-dfp-first-look-and-exchange-bidding/>

²⁰⁷ Index Exchange is an exchange that operates in the ad exchange market.

²⁰⁸ Vibram is a company that produces shoes.

winning at a clearing price of \$12. Vibram wins the impression through Index Exchange, which pockets 10% and passes on \$10.8 to DFP. DFP then takes 5% and passes on \$10.26 to the publisher as a payout.²⁰⁹ This example is illustrated in Figure 29 below.

Figure 29: Vibram wins an impression through Exchange Bidding because it clears the Enhanced Dynamic Allocation reserve set by an Altra direct deal



B. Impact of Header Bidding

150. Based on the exchange bidding definition above, along with the definitions of header bidding from Section III.B.2 and Dynamic Allocation and Enhanced Dynamic Allocation from Section IV, I can provide some results regarding the effects of header bidding. In my opinion, header bidding improves publisher revenue and fill rate²¹⁰ in comparison to the waterfall process. Furthermore, header bidding's auction mechanics generate higher revenue for publishers than Exchange Bidding's auction mechanics do.

²⁰⁹ This example abstracts away from details like ad buying tool fees, for clarity.

²¹⁰ The fill rate of the publisher is defined as the amount of impressions they were able to sell divided by the total amount of impressions received by the publisher. A higher fill rate implies a lower amount of impressions going unsold, given the total number of impressions.

1) Comparison of header bidding to the waterfalls with and without
(Enhanced) Dynamic Allocation

151. Header bidding increased revenue for publishers, in comparison to the waterfall.²¹¹ To illustrate this, consider the bid that might be submitted by an exchange. Under a waterfall approach, this exchange might not be called, in which case the highest bid from this exchange can be interpreted as 0. If called, the highest bid from this exchange would need to meet or exceed their personalized reserve in order to win the impression. In contrast, under header bidding

- a. The exchange is certainly called and in order to win must submit a bid that not only exceeds their reserve, but also exceeds the bids of other exchanges. Therefore, this exchange would submit a bid under header bidding that is at least as high as their bid under the waterfall process for all impressions.
- b. Furthermore, observe that the winning payment under header bidding is always the maximum bid that exceeds its reserve, whereas the winning payment under waterfalls is the first bid that exceeds its reserve. Therefore, even if individual exchanges' bids were the same in header bidding versus waterfalls, the publisher's revenue would go up because header bidding selects the highest bid instead of the first acceptable bid from the waterfall.
- c. Putting both claims together, header bidding would generate bids at least as high as waterfalls, and header bidding generates increased revenue from the same bids compared to waterfalls, and therefore generates increased revenue overall. In particular, the primary source of increased revenue is due to the simultaneous format of header bidding (as compared to the sequential format of the waterfall).

152. This is also confirmed by internal Google documents, one stating that "Pubs are making more money using [header bidding] (20%+ according to pubs) because they are able to get a per-query price from more demand sources."²¹²

²¹¹ To be clear, the two scenarios I consider are (a) all exchanges participate in the waterfall with some personalized reserves, and (b) all exchanges participate in header bidding with those same personalized reserves. If the publisher chooses to re-optimize reserves after switching to header bidding, their revenue could only increase further.

²¹² GOOG-DOJ-AT-02639830 at -41. April 2016. "Exchange Bidding (aka Jedi) LPS AMERICAS TRAINING."

153. The claims above hold even with the effect of (Enhanced) Dynamic Allocation.²¹³ I prove this claim in Appendix E. The proof follows from the description of Header Bidding in Section III.B.2 but requires case by case analysis. Again, the primary source of increased revenue is due to the simultaneous versus sequential format (even in light of AdX's Last Look advantage).

154. If header bidding enabled publishers to reach an expanded set of advertisers that they would not have otherwise reached through Google's ad tech stack, one would expect this additional competition to generate greater revenues for publishers.²¹⁴ In particular, one would expect existing advertisers to submit higher bids due to increased competition from additional advertisers. As long as existing advertisers' bids do not decrease in response to additional competition,²¹⁵ these additional advertisers themselves simply provide additional bids from which the publisher can take the maximum.

155. Google documents support this conclusion as well. An internal Google document²¹⁶ states that header bidding is advantageous for publishers because "they aren't getting all of the demand from AdX, so having other exchanges bid for inventory via HB increases Pub yield," showing that header bidding enables publishers to elicit bids from a wider range of advertisers compared to operating solely through the Google ad tech stack.

2) Comparison of header bidding to Exchange Bidding

156. The auction mechanics of header bidding (without Enhanced Dynamic Allocation) would generate increased revenue for publishers, as compared to all exchanges participating in Exchange Bidding. Without (Enhanced) Dynamic Allocation, header bidding is a clean first-price auction with personalized reserves and low take-rate.²¹⁷ Exchange Bidding is a clean first-price

²¹³ To be clear, the two scenarios I consider are: (a) non-Google exchanges participate in the waterfall with some personalized reserves and AdX has (Enhanced) Dynamic Allocation, and (b) non-Google exchanges participate in Header Bidding with those same personalized reserves and AdX has (Enhanced) Dynamic Allocation. In particular, for impressions that are eligible for a high priority line item, the two are equivalent. For impressions ineligible for a high priority line item, AdX's reserve in (a) is the maximum of all received header bids and all personalized reserves, and AdX's reserve in (b) is the maximum personalized reserve.

²¹⁴ While this may seem like an obvious claim, bidding behavior in first-price auctions is notoriously complex and counterintuitive phenomena are certainly possible. See Bernard Lebrun. "Existence of an Equilibrium in First Price Auctions." *Economic Theory* vol. 7, no. 3. 1996. pg. 421–443; Bernard Lebrun. "First Price Auctions in the Asymmetric N Bidder Case." *International Economic Review* vol. 40, no. 1. 1999. pg. 125–142; Bernard Lebrun. "Uniqueness of the equilibrium in first-price auctions." *Games and Economic Behavior* vol. 55, no. 1. 2006. pg. 131–151.

²¹⁵ Which would not happen in a second-price auction format, since it is truthful.

²¹⁶ GOOG-AT-MDL-001811992. June 2017. "Exchange Bidding / Platform StratOps Meeting."

²¹⁷ Client-side header bidding is free through Prebid, an open-source software. See Prebid. "Boost Programmatic Advertising Revenue." Accessed on June 3, 2024. <https://prebid.org/>

auction with personalized reserves (or non-personalized reserves, under UPR) where DFP collects a 5% take-rate on top of the ad exchange fee. Clearly, the former mechanics lead to increased revenue for publishers as compared to the latter.

C. Last Look

157. “Last Look advantage” is a phrase used to refer to AdX’s ability, under Dynamic Allocation and Enhanced Dynamic Allocation, to see the header bidding clearing price before submitting their own bid.²¹⁸ While the definition and mechanics of Dynamic Allocation do not change depending on whether other exchanges participated via the waterfall process or header bidding, the implications do. Here, I briefly elaborate on the related concepts.

158. I have previously described that Dynamic Allocation and Enhanced Dynamic Allocation offer a Last Look advantage to AdX when other exchanges participate in header bidding. AdX learns the highest header bidding bid before submitting its own. This is equivalent to a first-price auction where all bidders except for AdX submit their bids first, then AdX learns the highest submitted bid and submits its own.^{219, 220}

- 1) Last Look helps AdX clear impressions that would have otherwise been cleared by the header bidding winner, by paying a penny more

159. Last Look advantage likely helped AdX have a higher win rate in comparison to AdX’s win rate without Last Look advantage, by helping AdX clear impressions that would have otherwise been cleared by the header bidding winner.²²¹ I justify this claim below.

- a. I consider two cases: (a) all exchanges submit bids via a first-price auction with personalized reserves, and (b) all non-Google exchanges submit bids via a first-price auction with the same personalized reserves, and AdX’s Last Look

Server-side header bidding take rates vary by the service provider. For example, Ad Butler charges \$100 per month and \$0.001 per a thousand bids for its “Programmatic Advertising” service. AdButler. “Get the Industry-Leading Ad Server.” Accessed on May 31, 2024.

<https://web.archive.org/web/20231129061555/https://www.adbutler.com/pricing.html>

²¹⁸ And, for a period, the Exchange Bidding clearing price as well. GOOG-TEX-00000744 at -54. April 26, 2017. “Exchange Bidding (Jedi) Open Beta Sates Readiness Review.”

²¹⁹ I have analyzed an example in Section IV.A.2 to demonstrate how the Last Look Advantage helps AdX in this situation.

²²⁰ Even for sophisticated publishers, AdX learns some information about the highest submitted bid (and in particular, that the highest submitted bid is below its reserve price).

²²¹ This is my opinion in aggregate, after considering that some publishers used default options while others were sophisticated and increased Value CPMs of header bids to boost AdX’s reserve.

advantage lets it win instead as long as it submits a bid exceeding both its (same) personalized reserve and the highest first-price bid.²²²

- b. Let h denote the highest bid from a non-Google exchange that clears its reserve, v denote the highest value of any advertiser who bids with AdX, and r denote AdX's personalized reserve. There are a few cases to consider.
 - i. If either h or r exceeds v , then certainly AdX will not win in either auction format. Therefore, both formats have the same outcome.
 - ii. If v exceeds both h and r , then AdX certainly wins with Last Look advantage. This is because AdX's reserve will be the maximum of h and r , and its advertiser with value v will certainly submit a bid exceeding h and r ,²²³ and therefore some AdX bidder will win. If AdX instead participates with other exchanges in the first-price auction, AdX might lose. In particular, if AdX runs a second-price auction, their highest-value bidder will still bid v , but AdX would need a second highest bid (or reserve) exceeding h in order to win. Similarly, if AdX runs a first-price auction, their highest-value bidder would shade their bid, and without knowing h might shade their bid below h . Therefore, in any auction where AdX has a chance to win while participating in a first-price auction, AdX certainly wins with its Last Look advantage (and there are cases where AdX would win with Last Look but lose without it).
- c. As a result, the Last Look advantage helps AdX win more impressions that would have gone to the header bidding winner otherwise, assuming publishers and

²²² Recall that I have shown in Appendix D a natural example where optimal bidders would submit the same bids in both settings, although this would not hold in all settings. Recall also that this justification covers a default publisher who sets the same personalized reserves and does not boost header bids. I will briefly discuss how this analysis changes for sophisticated publishers.

²²³ This holds if AdX runs a second-price auction, because the highest-value advertiser will submit a bid of v . It also holds if AdX runs a first-price auction, because the highest-value advertiser will strictly prefer to submit a bid that has a chance of winning than a bid that will certainly lose (although it is possible that their bid will ultimately lose to a higher bid of another AdX bidder, some AdX bidder will certainly win).

advertisers who bid the same whether or not AdX has a Last Look advantage.²²⁴

²²⁵

160. In addition, when AdX wins these additional impressions, in many cases it pays just a penny more than the winning header bid and so does not increase publisher revenue. Specifically, in cases where AdX's highest value v exceeds the highest header bid h , but h exceeds both AdX's second highest value and its reserve, then AdX wins the impression with Last Look at price h (plus a penny), while without Last Look the header bid of h would have won. When some exchanges participate in header bidding while others participate in exchange bidding,²²⁶ only those who participate in header bidding are vulnerable to AdX's Last Look advantage.²²⁷ Specifically, the highest header bidding bid would become the reserve for AdX, whereas bids through Exchange Bidding are not revealed to AdX before AdX submits its own bid. That is, within AdX and Exchange Bidding exchanges, no one has a Last Look Advantage over the other, because their bids are submitted simultaneously.²²⁸ Exchanges that are integrated into Exchange Bidding see the same DFP reserve as AdX, hence they also have a Last Look advantage over header bidding exchanges.²²⁹

161. Therefore, one interpretation of Exchange Bidding is that it creates two tiers: Exchanges that participate in header bidding and exchanges that participate in Exchange Bidding together with AdX.²³⁰ Exchanges that participate in header bidding submit bids without seeing others' bids, and therefore have no Last Look advantage over anyone, and are vulnerable to AdX's and Exchange Bidding's Last Look advantage (placing them in the lowest tier).²³¹ AdX and Exchange Bidding exchanges have a Last Look advantage over exchanges that participate in header bidding, but do not have a Last Look advantage over each other, placing them in the top tier.

²²⁴ Recall again that I have shown a natural example in Appendix D where bidding identically in these situations is optimal for advertisers.

²²⁵ If a sophisticated publisher mildly inflates AdX's reserve specifically because of the Last Look advantage, or if advertisers bid similarly in these two cases, the same conclusions still qualitatively hold. If a sophisticated publisher drastically inflates AdX's reserve specifically because of the Last Look advantage, or advertisers drastically change their bids specifically due to AdX's Last Look advantage, the impact is less clear-cut and would require a complicated analysis weighing the benefits of Last Look versus the impact of an increased reserve and distinct bids.

²²⁶ Importantly, exchanges could participate in both. I am not claiming that this was the case, I am providing explanations based on a hypothetical.

²²⁷ Last Look was phased out in 2019, when AdX transitioned to a first-price auction format. See GOOG-TEX-00841386 at -88. "Adx First Price Auction." ("removing last look.")

²²⁸ See GOOG-TEX-00000744 at -54. April 26, 2017. "Exchange Bidding (Jed') Open Beta Sates Readiness Review." (diagram shows the '3p floor' entering AdX, but not other exchanges.)

²²⁹ GOOG-DOJ-AT-01815211 at -222. October 2019. "Open Bidding (fka Exchange Bidding)Training."

²³⁰ Again, exchanges can happen to be in multiple of these groups, since they can integrate into both header bidding and Exchange Bidding.

²³¹ As previously noted, if exchanges that are integrated into header bidding see the same DFP reserve as AdX, header bidders are also vulnerable to a Last Look from these exchanges.

Because a Last Look advantage is significant in first-price auctions, it would be natural for exchanges to want to remove the top tier's Last Look advantage over them and to gain a Last Look advantage over header bidders,²³² even though DFP takes a 5% fee on top of the clearing price when the winner is an Exchange Bidding exchange.^{233, 234}

VI. CONDUCT ANALYSIS: UNIFIED PRICING RULES

162. In this section, I provide an analysis of Google's Unified Pricing Rules (UPR), which was instated in 2019 (and in place today²³⁵) along with Google's ad exchange AdX's transition to the first-price auction format.²³⁶

163. I demonstrate that UPR leads to lower revenue for the publishers. I also demonstrate that UPR can lead to better win rate and revenue for Google's ad exchange AdX as well as for Google's ad buying tools and lower the win rate and revenue for rival exchanges and ad buying tools.

164. Prior to UPR, publishers could set different reserves that applied to different exchanges or different ad buying tools. Under UPR, publishers can no longer employ these personalized reserves to their full extent,²³⁷ because any reserve price set for non-guaranteed line items applies to all non-guaranteed line items.²³⁸ This reduces publisher choice by preventing them from setting personalized reserve prices. Publishers retain the ability to set personalized reserves on individual advertisers, but not on individual exchanges or ad buying tools.²³⁹ Furthermore, publishers may

²³² And after the change referenced in GOOG-DOJ-AT-01809483 went live, it would be further natural for exchanges to want a Last Look over header bidders. GOOG-DOJ-AT-01809483 at -89. March 2017. "Exchange Bidding in Dynamic Allocation (fka Project Jedi)."

²³³ Of course, the most natural auction format is to avoid creating tiers and a Last Look advantage at all, and to simply have all exchanges submit bids simultaneously without seeing each other's.

²³⁴ Last Look advantage was removed in 2019 during the implementation of Unified Pricing Rules and AdX's switch to the first-price auction format. See GOOG-TEX-00841386 at -89. "Adx First Price Auction."

²³⁵ Google. "Unified pricing rules." Accessed on May 31, 2024.

<https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en> (current Google Ad Manager documentation on UPR).

²³⁶ See generally Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. <https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/>

²³⁷ GOOG-AT-MDL-000875073 at -83. August 2019. "The Unified First Price Auction."

²³⁸ See Google. "Unified pricing rules." Accessed on May 31, 2024.

<https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en> (current Google Ad Manager documentation on UPR).

²³⁹ More specifically, publishers are free to set their reserve price at any level they desire, however this reserve price applies to all exchanges and ad buying tools.

set at most 200 total reserve prices at the advertiser level.^{240, 241} The excerpt in Figure 30 from an internal Google slide deck²⁴² discusses the reserve setting abilities publishers lose under UPR.

Figure 30: An excerpt from an internal Google document specifying that reserve prices under UPR applies to all non-guaranteed line items²⁴³

Summary of changes with Unified Pricing Rules		
	AdX Open Auction Pricing Rule	Unified Pricing Rule
Floor applies to	● Authorized Buyers (Ad Exchange)	<ul style="list-style-type: none"> ● Authorized Buyers (Ad Exchange) ● Exchanges on EBDA ● Non-guaranteed line items (excluding \$0 non-guaranteed and House)
Rules priority, overlapping rules	Floor from rule with higher priority will apply to the demand available	Maximum available floor will apply to the demand available
Branding Types	Branded, Semi-Transparent, Anonymous	Branded, Semi-Transparent (no different pricing per branding type)
Per- buyer floor	Yes	No
Blocks (buyer / advertiser)	Pricing Rule UI	Blocks migrate to Protections UI *

* Coming soon Google

165. As I outlined in Section II.B, there are several possible reasons why publishers might choose to set personalized reserve prices. First, they may wish to compensate for ad quality. A publisher may wish to charge more for low-quality ads (whose display negatively impacts user experience and indirectly causes the publisher financial loss) than high-quality ads (whose display

²⁴⁰ More precisely, any UPR reserve applies to all exchanges in Exchange Bidding as well as AdX. They also apply to line items that correspond to the header bidding winning bids. Note that publishers can configure any floors they like within their header bidding setup, including personalized reserves, but the header bidding line item will still be blocked if they are under the relevant UPR reserves. See generally Google, “Unified pricing rules.” Accessed on May 31, 2024. <https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en> (current Google Ad Manager documentation on UPR).

²⁴¹ At the initial announcement stage, Google stated that the system would allow 100 price floor rules. See Sarah Sluis, “Publishers Lash Out Against Google Over ‘Unified Pricing’ Changes” (April 18, 2019). Accessed on May 31, 2024. <https://web.archive.org/web/20221102060335/https:/www.adexchanger.com/online-advertising/publishers-lash-out-against-google-over-unified-pricing-changes/>; News Corp Australia, “Submission To the Australian Competition and Consumer Commission” (May 2020). Accessed on May 31, 2024.

[https://web.archive.org/web/20221012074940/http:/www.accc.gov.au/system/files/News%20Corp%20Australia%20\(1%20May%202020\).pdf](https://web.archive.org/web/20221012074940/http:/www.accc.gov.au/system/files/News%20Corp%20Australia%20(1%20May%202020).pdf) But based on the negative feedback, they increased the rule limit to 200. See GOOG-TEX-00594205 at -11. “The Unified First Price Auction Best Practices.” (“Based on your feedback, we are improving Unified Pricing Rules; Coming soon; Unified pricing rules limit increased to 200.”)

²⁴² GOOG-AT-MDL-000875073. August 2019. “The Unified First Price Auction.”

²⁴³ GOOG-AT-MDL-000875073 at -83. August 2019. “The Unified First Price Auction.”

might be neutral, or at least less negative, to user experience). If a publisher believes that ads served through AdX on average cause more financial loss in comparison to ads served through other exchanges, it makes sense to set a higher personalized reserve to AdX than to other exchanges. Furthermore, publishers can block ads coming from an exchange altogether if they deem the ads particularly damaging, by setting a very high reserve for that exchange. Personalized reserve prices allow the publisher to express their financial preference for different ads and improve efficiency. Google is also aware of the publishers' attention to ad quality, stating that "many AdX publishers are very sensitive to ads that they feel reflect badly on their brand and unwilling to risk any exposure at all; when such ads [...] appear on their pages, they may react by pulling their inventory from AdX altogether."²⁴⁴

166. Another reason why publishers may wish to set personalized reserves is that they might want to charge more to an exchange that tends to produce higher bids since these exchanges will generate higher revenue for the publishers. One exchange might tend to produce higher bids than another for two key reasons: (a) the exchange simply has access to a larger number of advertisers²⁴⁵ or (b) the exchange has behavior that causes them to submit higher bids with access to a comparable advertiser number (such as Google's Project Bernanke, which I analyze in Section VIII).

1) UPR can lead to inefficient outcomes

167. Imagine a publisher finds that ads served through AdX to on average cause financial loss of \$5 more than ads through OpenX do, and therefore chooses to set a personalized reserve \$5 higher on AdX (to have concrete numbers in mind, say that OpenX causes \$2 loss on average, while AdX causes \$7 on average).²⁴⁶ Hence the publisher sets a reserve of \$15 for AdX and \$10 for OpenX. AdX submits a bid of \$13, and OpenX submits a bid of \$11. Then with personalized reserves, OpenX will win the impression, and this is the efficient outcome (yielding payoff \$9 for the publisher; \$11 of revenue minus \$2 of harm). Under UPR, the publisher instead must set the same reserve on both exchanges. That reserve might be less than or equal to \$13 which will result in AdX winning the impression, which nets the publisher more revenue. But the publisher's


²⁴⁴ GOOG-DOJ-15769995 at -5. May 2017. "Protecting Publishers from Objectionable Ads - Proposal."

²⁴⁵ And this itself might be due to 'good business' of providing a better service, or due to some of the other conduct allegations in this case where Google products preference AdX on both the buy-side (causing advertisers to end up using AdX even though another exchange might provide a better price or product) and the sell-side (causing the same bids through AdX to have a better shot at winning) of the market.

²⁴⁶ Note that setting optimal reserves is a complex problem and setting a reserve that is exactly \$5 higher on AdX because ads cause exactly \$5 more harm on average is not likely to be optimal. However, it is a sensible heuristic, and the key takeaways from the example are not driven by the particular choice of personalized reserve.

true gain is \$6; \$13 of revenue minus \$7 of harm. The publisher might instead set a single reserve that is higher than \$13 which will result in no sale. In this case, there is no single reserve price which results in the efficient outcome of OpenX winning the impression, because UPR removes the ability of the publisher to set personalized reserve prices. This example is illustrated in Figure 31 below.

Figure 31: The efficient outcome cannot be achieved in this auction due to UPR



168. Optimal reserves improve revenue by putting pressure on bidders to increase their bids.²⁴⁷ Imagine that the publisher knows that the top bid from AdX tends to be between \$12 and \$16, whereas the top bid from OpenX tends to be between \$10 and \$12, and for simplicity of this example assume that each exchange has a single advertiser.²⁴⁸ The publisher might therefore set a reserve of \$13 on AdX and \$10 on OpenX. Under these reserves, the AdX bidder would optimally shade any bids above \$13 down to \$13, and the OpenX bidder would optimally shade all bids down to \$10. When AdX's bidder has a value of \$15, these personalized reserves and

²⁴⁷ Because UPR was implemented as AdX switched to a first-price auction format, and during a time when the majority of the intermediated display advertising ecosystem was also using first-price auctions. See Sarah Sluis. "Google Switches To First-Price Auction" (March 6, 2019). Accessed on May 31, 2024.

<https://web.archive.org/web/20220910040643/https://www.adexchanger.com/online-advertising/google-switches-to-first-price-auction/> ("Google Ad Manager will be the last major exchange to switch to first-price auctions. Other exchanges tested or rolled out first-price auctions starting in 2017.")

I focus my analysis on this case.

²⁴⁸ To make this example mathematically rigorous, imagine that there is a single AdX bidder whose value is distributed uniformly on [12,16], and there is a single OpenX bidder whose value is distributed uniformly on [10,12]. In this case, the revenue-optimal personalized reserves are \$13 on AdX and \$10 on OpenX. Moreover, the example bids for AdX and OpenX are optimal in this case.

optimal advertiser behavior results in a revenue of \$13. When AdX's advertiser instead has a value of \$12, these personalized reserves and optimal advertiser behavior results in a revenue of \$10, even in cases where AdX's advertiser does not clear the reserve price. But under UPR the publisher must set the same reserve for both exchanges. If that reserve exceeds \$12, then it likely excludes OpenX,²⁴⁹ and acts as a reserve on AdX alone. In particular, if AdX fails to meet the reserve, the publisher likely foregoes the fallback option of \$10 they were able to achieve with personalized reserves. If instead that reserve is below \$12, it puts no pressure on AdX's advertiser, which would optimally shade its bid down to \$12.

169. In summary, personalized reserves allow the publisher to both set a competitive reserve for higher-priced exchanges and still collect some revenues when higher-priced exchanges fail to meet those reserves, while UPR forces the publisher to choose between one or the other. Furthermore, personalized reserves allow the publishers to screen for ad quality, by either charging more to compensate for the low quality or blocking the low quality ads altogether. UPR disables publishers from effectively using reserve prices to screen for ad quality.

B. Impact of Unified Pricing Rules on Publishers

170. In this subsection, I explain why UPR leads to a revenue loss for publishers. Under UPR, publishers lose the ability to maximize their revenues because they cannot set personalized reserves for exchanges and ad buying tools.

1) UPR prevents publishers from maximizing their revenues

171. The ability to set personalized reserves for each exchange is an important revenue-optimization tool for publishers. UPR took away this ability, as discussed above. This is true when publishers are interacting with an advertiser for the first time, as well as with an advertiser they have sufficient information about.

- a. For new advertisers, or advertisers with whom the publisher has not interacted enough to be able to set a suitable reserve price, publishers would rely on coarse information to set appropriate reserves. For example, advertisers that choose to transact through AdX may tend to be materially different (both in willingness to pay and in ad quality) than advertisers who chose to transact through non-Google exchanges. In order to maximize revenue, publishers would still want to set as

²⁴⁹ And in the mathematically rigorous example, completely excludes OpenX.

good a reserve as possible on new advertisers for whom they have not yet decided on an appropriate advertiser-specific reserve. The natural mechanism by which to do this is to set a single personalized reserve on AdX (or the non-Google exchanges) that applies to all advertisers, including new ones, who transact through AdX. Note that the ability to set personalized reserves to advertisers does not address this source of revenue loss since per-advertiser reserves are only meaningful with sufficient information about the advertiser, whereas a per-exchange reserves are necessary to maximize revenues based on available information on new advertisers.

- b. Even when the advertiser is known to the publisher, publisher revenue can depend on the exchange through which that advertiser's bid occurred. The same advertiser's initial bid is processed differently through AdX than non-Google exchanges.²⁵⁰ Therefore, a revenue-maximizing publisher would naturally want to set different reserves for exchanges that engage in different conducts, even if they reach the same advertiser pool.
- c. Finally, even when the advertiser is known to the publisher, *and* that advertiser is known to transact through exactly one exchange, the '200 Rules' limit that comes with UPR would still lead to revenue loss among publishers because it prevents the implementation of revenue-maximizing strategies that use personalized reserves. When there are many advertisers potentially interested in an impression, revenue-maximization requires a granular reserve for each advertiser, as explained above. Limiting the number of allowed reserve prices to 200 limits the publisher's ability to maximize revenues.

2) UPR can lead to lower quality ads for the publishers

172. As discussed above, UPR can decrease the ad quality that the publishers face. This is because publishers are not able to filter for ad quality across different exchanges by employing personalized reserves. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

²⁵⁰ For example, GDN's Project Bernanke inflates the highest bid sent to AdX. Therefore, the same advertiser's bid is likely higher after being processed through AdX than a non-Google exchange.

[REDACTED]
[REDACTED] 251

173. In sum, UPR would lead to a revenue loss for publishers. UPR would decrease publisher welfare even further by preventing them from effectively screening for low quality ads.

C. Impact of Unified Pricing Rules on Exchanges and Ad Buying Tools

174. In this subsection, I demonstrate that UPR can lead to a better win rate and higher revenue for Google's ad exchange AdX as well as Google's ad buying tools, and lower win rate and revenue for rival exchanges and ad buying tools. I provide explanations as to why these are plausible outcomes of UPR. I show that if AdX was facing high reserves compared to other exchanges compared, then it is plausible to expect that AdX revenue increased as a result of UPR. Furthermore, I explain that UPR can hurt exchanges and ad buying tools that facilitate the transaction of high quality ads.

175. Under UPR, for each advertiser, AdX receives a reserve price that is no higher than any other non-guaranteed line item, including other exchanges. This is true for exchanges that are called via Exchange Bidding, the waterfall²⁵² or header bidding.²⁵³ In other words, for a given impression, all exchanges face the same reserve. This can benefit AdX if it previously faced higher reserves compared to other exchanges.

176. Since setting personalized reserves is a complex problem, it is likely that publishers use general heuristics.²⁵⁴ One plausible such heuristic is the **Treat-As-Single Heuristic**, where publishers set personalized reserves for each individual exchange as if it were the only exchange (optimally according to this assumption, using Myerson's (1981)²⁵⁵ reserve),²⁵⁶ and publishers set

²⁵¹ GOOG-AT-MDL-004017152. ("[REDACTED].")

²⁵² Note that exchanges called through Exchange Bidding must pay the Exchange Bidding fee, and that exchanges that participate through the waterfall may not be called at all.

²⁵³ Exchanges that are called through header bidding face a reserve price at least as high as the UPR reserves but may face a higher personalized reserve. This is because UPR applies to all non-guaranteed line items, which include header bidding line items, but the header bidding technology still allows for personalized reserves, so publishers can still choose to set personalized reserves in the header bidding auction, with the condition that these must be higher than the applicable UPR reserve set in DFP.

²⁵⁴ A heuristic for a publisher in this case can be thought of a guideline that determines their actions.

²⁵⁵ Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73. Subsequent works are also relevant.

²⁵⁶ The personalized reserves set under the Treat-As-Single Heuristic are a commonly studied heuristic in auction theory under the name "VCG with Monopoly Reserve." See, e.g., Hartline and Roughgarden. "Simple versus Optimal Mechanisms." *Proceedings of the 10th ACM Conference on Electronic Commerce*. 2009. pg. 225-234.

UPR reserves by setting Myerson's optimal reserve for a hypothetical "average exchange."²⁵⁷ Another plausible heuristic is the **eCPM Heuristic**, where publishers set personalized reserves for each individual exchange equal to the exchange's historical average bid on similar impressions (after adjusting for value), and publishers set UPR reserves equal to the average historical bid on similar impressions (across all exchanges, after adjusting for value).^{258 259}

177. If publishers use the Treat-As-Single Heuristic, then whatever exchange faces the highest personalized reserve would face a lower reserve under UPR. I provide a proof of this claim in Appendix F. The proof establishes that the Treat-As-Single Heuristic sets a UPR reserve that is somewhere between the minimum and maximum personalized reserves, under a standard 'regularity assumption' on the bids typically produced by each individual exchange. Therefore, whichever exchange faces the highest reserve will also face a lower UPR reserve, and whichever exchange faces the lowest reserve will also face a higher UPR reserve. If publishers use the eCPM heuristic, then whatever exchange faces the highest personalized reserve would also face a lower reserve under UPR. The argument again establishes that the eCPM heuristic sets a UPR reserve somewhere between the minimum and maximum personalized reserves (and does not require a regularity assumption). Therefore, again, whichever exchange faces the highest personalized reserve will also face a lower UPR reserve, and whichever exchange faces the lowest reserve will face a higher UPR reserve. While it is too complex to guess precisely how publishers set reserves, these two natural heuristics demonstrate similar behavior and result in UPR reserves somewhere between the minimum and maximum personalized reserves, which implies lower UPR reserves for exchanges that received the highest personalized reserves. It is therefore natural to expect other publisher heuristics to follow a similar analysis.

- 1) If AdX faced the highest personalized reserve pre-UPR, then AdX would transact more impressions and have increased revenue under UPR

178. If AdX faced the highest personalized reserve pre-UPR, then AdX would likely transact more impressions and have increased revenue under UPR. There are two steps to understand why this is the case. First, if AdX faced the highest personalized reserve pre-UPR, then it is natural to expect AdX's UPR reserve to be lower, as I explained previously. Second, I now argue that

²⁵⁷ To be clear, by "average exchange," I mean that the publisher considers an exchange with value distribution whose cumulative distribution function is the average of the cumulative distribution functions of all individual exchanges. In more formal terms, if exchange i has cumulative distribution function F_i , then the "average exchange" of n exchanges has the cumulative distribution function $\sum_{i=1}^n F_i/n$.

²⁵⁸ To be clear, by "adjusting for value" I mean the following: "If the harm caused by displaying ads from this exchange is on average \$X CPM, set the reserve for this exchange equal to its eCPM plus \$X."

²⁵⁹ I give more detailed definitions for these heuristics in Appendix F.

AdX would transact more impressions and have increased revenues under UPR, assuming that indeed AdX's UPR reserve is lower than its pre-UPR personalized reserves. Specifically, I now compare the win-rate and competition faced by a particular exchange in the following two scenarios: (a) Personalized reserves, with this particular exchange's being the highest, and (b) unified reserves, and less than this particular exchange's personalized reserve in (a). I conclude that, if the bidder behavior results in efficient outcomes,²⁶⁰ this particular exchange:

- a. Wins more impressions under unified reserves. More specifically, for each impression that this particular exchange would have won without UPR, this particular exchange also wins under UPR. In addition, there are impressions this particular exchange would not have won without UPR that this particular exchange now wins under UPR.
- b. Faces less competition faced for impressions it wins on average under unified reserves. That is, for the set of impressions that this particular exchange would have won without UPR, this particular exchange also wins these impressions under UPR. Moreover, the competition faced on these impressions is no higher under UPR than without UPR.²⁶¹

I elaborate on these points further in Appendix F.

179. Any exchange whose pre-UPR personalized reserve is lower than their UPR reserve would experience a lower yield and lower revenues under UPR, in comparison to the counterfactual of not imposing UPR on publishers. Moreover, if AdX faced the highest personalized reserve pre-UPR, some other exchanges must face a higher UPR reserve than their pre-UPR personalized reserves under either of the two natural heuristics, and therefore would achieve a lower yield and revenues under UPR in comparison to the counterfactual of not imposing UPR on publishers. I compare the outcomes for exchange i under two sets of reserves: (a) Personalized, where exchange i faces a reserve of r_i , and (b) a unified reserve of $r \geq r_i$. I conclude that, if the bidder behavior results in efficient outcomes, Bidder i would:

²⁶⁰ The outcome of an auction is efficient if and only if, among the bidders who submitted a bid higher than their personalized reserves, the bidder with the highest valuation of the auctioned item wins the item.

²⁶¹ Formally, what I mean by this is that treating the exchange as a single bidder, that exchange *could* pay less on average and still win the same impressions under UPR as pre-UPR. The exchange may still choose to pay the same or more (i.e., by setting a reserve to its advertisers that exceeds the competition faced, using a program like Google's Reserve Price Optimization). Or the exchange may instead choose to charge its advertisers the same as when it would win pre-UPR, and simply pass on less payout to the publisher (i.e., using a program like Google's Dynamic Revenue Sharing, although that family of conducts never occurred at the same time as UPR).

- a. Win fewer impressions under UPR. That is, for each impression that exchange i would have won under UPR reserve of r , exchange i certainly wins under personalized reserves of r_i .
- b. Faces higher competition for the impressions they win on average. That is, for the set of impressions that exchange i wins under UPR at reserve r , exchange i also wins these impressions at personalized reserves r_i . Moreover, the competition faced on these impressions is no higher under personalized reserves r_i than under UPR at reserve r .²⁶²

I elaborate further on these points in Appendix F.

- 2) UPR would decrease the win rates and revenues of exchanges and ad buying tools that typically transact high quality ads

180. UPR would decrease the win rates and revenues of exchanges and ad buying tools that typically transact high quality ads. In general, UPR requires that bids are treated identically from all ad buying tools and all exchanges. In contrast, personalized reserves allow publishers to preference bids of the same dollar value from ad buying tools and exchanges that transact higher quality ads. Under personalized reserves, it is plausible that publishers set higher personalized reserves for ad buying tools and exchanges that facilitated transactions of lower quality ads in comparison to others, which might enable, for example, a bid of \$5 from a high-quality ad to be selected over a bid of \$5.01 from a low-quality ad. But, under UPR, all bids are treated identically and a \$5.01 bid from a low-quality ad will always be selected over a \$5 bid from a high-quality ad.²⁶³ Therefore, exchanges that typically transact high-quality ads would have lower win rates and revenues under UPR.²⁶⁴

181. For example, if the typical values of bidders are identical across exchanges, but the typical ad quality differs, the exchange with the highest ad quality would face a lower reserve with personalized reserves compared to UPR, under natural heuristics such as the Treat-As-Single

²⁶² Again, what I formally mean by this is that treating the exchange as a single bidder, that exchange *must* pay more on average to win any impression that it wins under UPR as compared to what it must pay pre-UPR. As noted previously, the exchange can always choose to set a higher reserve on its advertisers than it must pay, so all exchanges prefer to be less constrained in what they *must* pay to a publisher in order to win an impression.

²⁶³ The lone exception to this is if the \$5.01 advertiser is subject to one of the 200 rules. As previously noted, 200 rules may be sufficient to screen out some undesirable advertisers but is insufficient to express a general preference for high-quality over low-quality ads in the same manner as personalized reserves.

²⁶⁴ Due to the complexity in optimizing reserves and equilibria of first-price auctions, it is challenging to make strong predictions about what might happen in practice from an auction theory perspective alone. The analysis states what I would see as the most likely outcome in such a scenario, based on my expertise in auction theory.

Heuristic or the eCPM heuristic. Exchanges with high ad quality would therefore suffer lower win rates and revenues under UPR. Between two exchanges with identical typical values but distinct typical ad quality, the exchange with higher ad quality would face a lower Myerson reserve when treated as the only exchange, and this would result in higher win rates and higher revenues under the Treat-As-Single Heuristic. Similarly, between two exchanges with identical typical values but distinct typical ad quality, the higher ad quality would have a stronger value-adjusted eCPM (because the eCPMs are identical, but value-adjusting works more strongly in favor for higher-quality ads). I provide a more formal explanation for this result in Appendix F.

182. In sum, assuming that AdX faced the highest reserve pre-UPR, UPR would naturally benefit Google's ad exchange AdX both in win rate and in revenue.²⁶⁵ Furthermore, it negatively affects some non-Google exchanges and ad buying tools, especially those that tend to transact high quality ads.

VII. CONDUCT ANALYSIS: DYNAMIC REVENUE SHARING

183. In this section, I analyze Google's Dynamic Revenue Sharing (DRS)²⁶⁶ conduct and its variants, which went on from 2015 to 2019. I find that:

- a. Dynamic Revenue Sharing version 1 (DRSv1) increased AdX win rate and revenue and decreased non-AdX exchanges' win rates and revenues, compared to no DRS,
- b. Dynamic Revenue Sharing version 2 (DRSv2), in comparison to both no DRS and DRSv1, decreased advertiser payoff, increased AdX win rate and revenue, decreased non-AdX exchange's win rates and revenues, and may also decrease publisher revenue.
- c. Google concealed information that is vital to advertisers and important to publishers by concealing DRSv1 from them. At least some of Google's communication regarding DRSv2 was misleading.

²⁶⁵ If AdX faced a relatively high reserve pre-UPR, but not the *highest* reserve pre-UPR, all of my conclusions qualitatively hold for the same reasons. I chose to present the case where AdX faces the *highest* reserve in the interest of crisp statements and clean proofs.

²⁶⁶ This conduct is sometimes referred to as "sell-side DRS." GOOG-AT-MDL-003849201. Email thread, "Subject: Re: Partner revshare scaling factor alerts." ("The timing seems to match with the sellside DRS experiment.")

184. Before the DRS program, Google imposed a contractually mandated²⁶⁷ per-auction take rate that was typically equal to the 20% of what is submitted by the ad buying tool.²⁶⁸ This implies that Google either (a) took 20% of the second highest net bid²⁶⁹ and passed on the remaining 80% to the publisher if the remaining 80% of the second highest net bid was higher than the publisher's reserve or (b) paid the publisher their reserve and charged the advertiser an amount that corresponds to 125% of the reserve,²⁷⁰ since the 80% of that amount would exactly be the reserve.

185. As a result, prior to DRS, AdX ran the following auction:

- a. AdX learns its reserve r from the ad server. Due to the static take rate of 20% imposed by AdX, the highest bid submitted to the AdX auction must be at least $r/0.8$ to win the impression.²⁷¹
- b. AdX solicits bids from the ad buying tools. Let b_1 denote the highest received bid and b_2 denote the second highest.²⁷²
- c. If the second highest bid is above the effective reserve²⁷³ (*i.e.*, $b_2 \geq r/0.8$), then the clearing price is b_2 , AdX charges $0.2 * b_2$ as its fee, and passes on $0.8 * b_2$ to the publisher.
- d. If the effective reserve is in between the highest and second highest bids (*i.e.*, $b_1 \geq r/0.8 > b_2$), then the clearing price is $r/0.8$, AdX charges $0.2 * r/0.8 = r/4$ as its fee and passes on $0.8 * r/0.8 = r$ to the publisher.
- e. If the highest bid is below the effective reserve (*i.e.*, $r/0.8 > b_1$), then the highest bid fails to clear the reserve price after AdX takes its rate, hence AdX fails to win the impression.

²⁶⁷ [REDACTED]

[REDACTED].")

²⁶⁸ [REDACTED]

[REDACTED].")

²⁶⁹ That is, the bid submitted by the ad buying tool, after having already taken the ad buying tool fee.

²⁷⁰ Abstracting away from the ad buying tool fee.

²⁷¹ AdX's take rate is taken based on the clearing price.

²⁷² For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

²⁷³ By "effective reserve price," I mean the reserve price plus the take rate AdX take if the reserve price ends up being the clearing price.

186. This is a regular second-price auction in which bids are adjusted by the exchange's take rate. Until 2019, Google stated that it ran a second-price auction.²⁷⁴ The second price auction is truthful, which means that it is in the advertiser's best interest to bid their true value for the impression. This has many desirable properties, as noted in Section II. I note that the first price auction is not truthful.²⁷⁵

187. In order to analyze the impacts of the different types of DRS conducts, I first introduce a **dirty second-price auction**, as defined in Google's internal documentation.²⁷⁶ A dirty second-price auction has two parameters, a hard floor r , and a soft floor $q \geq r$. There are many ways to implement a dirty second-price auction, but they will all have the following properties when only the highest bid b_1 exceeds the hard floor r .

- a. If the highest bid is above the soft floor (*i.e.*, $b_1 \geq q$), then the highest bidder wins and pays q . This is the same as the regular second price auction with reserve q .
- b. If the highest bid is below the hard floor (*i.e.*, $b_1 < r$), then no one wins. This is same as the regular second price auction with reserve q .
- c. If the highest bid is in between the soft and hard floors (*i.e.*, $q > b_1 \geq r$), then the highest bidder wins and pays b_1 . This differs from the regular second price auction with any reserve, due to the fact that the highest bidder is paying their own bid. Hence for this specific case, the dirty second price auction gives the same outcome as a first price auction.

188. Notice that with dirty second price auctions, under some conditions, the highest bidder pays their own bid. This is the same as the first price auction. Since the first price auction is not truthful, this means that the dirty second price auction is not truthful.²⁷⁷ To see this, suppose there is a bidder who has a value v that exceeds r but not q , and that no other bidder has a value exceeding r . If this bidder bids their true value, they will win the impression and pay v . If instead

²⁷⁴ See Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. <https://web.archive.org/web/20240122142404/https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/>

²⁷⁵ See Section II for more information about truthful auctions.

²⁷⁶ GOOG-NE-05279363 at -73. "Bidding in adversarial auctions."

²⁷⁷ Under a dirty second price auction, the bidders have an incentive not to be truthful only under some conditions, but a truthful auction format is defined as always incentivizing the bidders to bid their true values. Since the dirty second price auction fails to do that under some conditions, it is fails to be truthful in general.

they submit a bid, b , between the reserve, r , and their value, v , they will win the impression and pay $b < v$ which is strictly better than bidding their true value.

A. Dynamic Revenue Sharing v1

189. DRS was launched in August 2015²⁷⁸ without announcing it to publishers or advertisers.²⁷⁹ In its first iteration, DRSv1, AdX dynamically decreased its take rate to be lower than 20% to win impressions that it would not have if the take rate was kept at 20%. AdX decided to either impose a 20% take rate or decrease the take rate depending on a few factors including comparisons between the first and the second highest bid and the publisher reserve,²⁸⁰ as well as the average take rate among the auctions for that publisher's impressions in that billing period.²⁸¹

1

²⁷⁸ GOOG-TEX-00777528 at -30. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)." (The email from September 2nd, 2015 states that "Last week we launched Dynamic sell-side Revenue Share (DRS).")

279

²⁸⁰ GOOG-NE-06864639 at -43, May 9, 2014. "Dynamic Sell-side Revshare on AdX." (under the subsection "Cases.")

281

²⁸² This conduct occurred exclusively during a period where AdX ran a second-price auction.

²⁸³ For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

191. DRSv1 impacts the auction only in one case where the highest bid is high enough to clear the reserve price, but not high enough to do so if AdX takes its full fee of 20% of the clearing price. When that happens, AdX dynamically decides to decrease its take rate so that it returns a successful bid to the ad server and wins the impression.

192. To illustrate how DRSv1 works, imagine an impression arrives, and the publisher reports a price floor of \$10.²⁸⁸ AdX solicits bids and receives top-two bids of \$20 and \$10. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is \$12.5, because $\$20 \geq \$12.5 > \$10$. AdX takes a 20% take rate of \$2.5 and passes on \$10 to the publisher and DRSv1 has no impact. This example is illustrated in Figure 32 below.

²⁸⁸ For ease of exposition, I use ‘price floor’ for when the publisher sets a reserve on an exchange, and ‘reserve’ for when an exchange sets a reserve on advertisers.


Figure 32: An auction where DRSv1 has no impact



193. Imagine another impression arrives, and the publisher again reports a price floor of \$10. AdX solicits bids and this time receives top-two bids of \$9 and \$8. In this case, with both a regular second-price auction and DRSv1, the impression is not transacted with AdX, because even with a take rate of 0% AdX still cannot clear the reserve.²⁸⁹ This example is illustrated in Figure 33 below.

²⁸⁹ This is a key distinction between DRS and Project Bernanke. Project Bernanke in some sense is a 'step further' than DRS because it further tries to transact at a negative margin. See Section VIII for further details on Project Bernanke.

Figure 33: An impression that AdX could not clear even with DRSv1




194. Imagine a final impression arrives, and the publisher again reports a price floor of \$10. AdX solicits bids and this time receives top-two bids of \$11 and \$10. Under a regular second-price auction with a take rate of 20%, the impression would not be transacted with AdX, because that would require a top bid of at least \$12.5 for a 20% take rate to still result in a payment above \$10. Under DRSv1, provided that DRSv1 is not throttled,²⁹⁰ DRSv1 would instead clear the transaction, charge the winner \$11, collect a \$1 fee for a take-rate of 9.1%. and pass on \$10 to the publisher. Note that this lowers the average take-rate across the billing period.²⁹¹ This example is illustrated in Figure 34 below.

²⁹⁰ If DRSv1 is throttled, DRSv1 will just execute as a regular second-price auction.

²⁹¹ In the final example, it is important to note that the advertiser is paying their bid. In particular, if the advertiser had reported a bid of \$12 instead, they would have been charged \$12 (and AdX would have taken \$2, for a take-rate of approximately 16.66%). If the advertiser had reported a bid of \$10.1, they would have been charged \$10.1 (and AdX would have taken \$0.1 CPM, for a take-rate of approximately 1%). That is, while DRSv1 acts like a truthful second-price auction for some bids, it is instead a first-price auction in the 'dynamic range' from \$10 to \$12.5, where AdX will charge the winner their bid.

Figure 34: An impression cleared by AdX with the DRSv1 dynamic adjustment of the take rate



195. DRSv1 is a dirty second-price auction and as a result, DRSv1 is not truthful. Specifically, whenever DRSv1 charges the highest bidder their bid and pays the reserve r to the publisher, it is a dirty second-price auction with a hard reserve of r and a soft reserve equal to $r/0.8$. This is because when the highest bid is the unique bid above the hard reserve r and below the soft reserve $r/0.8$, they win and pay their bid.

B. Dynamic Revenue Sharing v2

196. With DRSv2, AdX dynamically adjusted its take rate to sometimes be higher or lower than 20% (while maintaining the average take rate at 20% across auctions cleared by AdX) to win impressions that it would not have if the take rate was kept at 20% on a per-query basis.²⁹² In contrast, under DRSv1, AdX only adjusted take rates to be below 20%.²⁹³ AdX decided to increase or decrease the take rate from 20% depending on a few factors such as the comparisons

²⁹² GOOG-NE-13207241 at -1. "AdX Dynamic Revshare v2: Launch Doc." [REDACTED]

²⁹³ GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)." ("We lower the revenue share per query as needed.")

between the first and the second highest bid and the publisher reserve, as well as debt balances for the publishers and advertisers.²⁹⁴

197. DRSv2 was launched in the second half of 2016.²⁹⁵ Google announced DRSv2 when it was launched.²⁹⁶ The publishers were allowed to opt out of DRSv2, however, if they did, Google turned off DRSv1 for these publishers as well.²⁹⁷ Advertisers and ad buying tools could not opt out of DRSv2.²⁹⁸

[REDACTED]

[REDACTED]

[REDACTED]

²⁹⁵ GOOG-NE-04934281 at -81. July 30, 2018. "Dynamic Revenue Share." (describing release dates of "feature flags" (functionality), including that it launched into AdX UI in June 2016 and came into effect in August 2016.)

²⁹⁶ GOOG-NE-06842715 at -20. May 10, 2016. "AdX Auction Optimizations." (describing that DRS would be announced in June 2016.)

²⁹⁷ GOOG-NE-04934281 at -86. July 30, 2018. "Dynamic Revenue Share." ("You may choose to opt-out of revenue share based optimizations in the AdX UI. If you opt-out we will apply your contracted revenue share to every Open Auction query and you will not benefit from the increased revenue from this optimization.")

²⁹⁸ GOOG-NE-04934281 at -85. July 30, 2018. "Dynamic Revenue Share." ("Q: Can buyers opt out? A: Revenue share based optimizations are controlled by sellers only.")

²⁹⁹ For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

[REDACTED]

[REDACTED]

[REDACTED]


[REDACTED]

199. Under DRSv2, AdX is allowed to either increase the take rate on a per impression basis or decrease it, as it can be seen from its description. This is one important difference in comparison to DRSv1 where AdX was only allowed to decrease the take rate. AdX decides to increase or decrease the take rate based on the average take rate it applied to that publisher in the billing period. If it is close enough to the contractual requirement of 20%, then AdX sometimes decreases the take rate to win impressions that it would not have won otherwise. If it is much lower than 20% and the top bid is high enough compared to the floor, AdX increases the take rate to recoup the losses it incurred in auctions where it decreased its take rate.

200. To illustrate how DRSv2 works, imagine an impression arrives, and the publisher reports a price floor of \$10. AdX solicits bids and receives top two bids of \$20 and \$15. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is \$15, because \$15 is greater than $\$10/0.8 = \12.5 . AdX takes a 20% take-rate of \$3 and passes on \$12 to the publisher. Under DRSv2, if the winning advertiser has some debt from prior impressions, DRSv2 then

additionally charges the advertiser \$0.5, clears \$0.5 of its debt, keeps a 20% take rate on this and passes \$0.4 to the publisher. In addition, if the publisher has some debt from prior impressions, DRSv2 then additionally withholds \$0.2 from the publisher and clears \$0.2 of its debt. So altogether, the advertiser pays \$15.5 instead of \$15, and the publisher is paid \$11.4 instead of \$12. This example is illustrated in Figure 35 below.

Figure 35: An impression is cleared by AdX with the DRSv2 dynamic adjustment of the take rate. Advertiser pays more, publisher earns less due to debt clearing



201. Imagine that another impression arrives, and the publisher again reports a price floor of \$10. AdX solicits bids and this time receives top-two bids of \$20 and \$10. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is \$12.5, because \$20 is greater than $\$12.5 > \10 . AdX takes a 20% take rate of \$2.5 and passes on \$10 to the publisher. Under DRSv2, if the winning advertiser has some debt from prior impressions, DRSv2 then additionally charges the advertiser \$0.5, clear \$0.5 of its debt, keeps a 20% take rate on this and passes \$0.4 on to the publisher. In addition, if the publisher has some debt from prior impressions, DRSv2 then additionally withholds \$0.2 from the publisher and clears \$0.2 of its debt. So altogether, the advertiser pays \$13 instead of \$12.5, and the publisher is paid \$10.2 instead of \$10. This example is illustrated in Figure 36 below.

Figure 36: An impression is cleared by AdX with the DRSv2 dynamic adjustment of the take rate. Advertiser pays more, publisher earns more due to debt clearing



202. Imagine another impression arrives, and the publisher again reports a price floor of \$10. AdX solicits bids and this time receives top two bids of \$9 and \$8. In this case, with both a regular second-price auction, DRSv1, or DRSv2, the impression is not transacted with AdX,³⁰³ because even with a take-rate of 0% AdX still cannot clear the price floor. This example is illustrated in Figure 37 below.

³⁰³ Note that the impression might be sold to another exchange or go unsold.

Figure 37: An impression that AdX could not clear even with the DRSv2 dynamic adjustment of the take rate




203. Imagine that a final new impression arrives, and the publisher again reports a price floor of \$10. AdX solicits bids and this time receives top two bids of \$11 and \$10. In this case, with a regular second-price auction and a static take-rate of 20%, the impression would not be transacted with AdX.³⁰⁴ Because that would require a top bid of at least \$12.5 for a 20% take-rate to still result in a payment above the price floor of \$10. With DRSv1 or DRSv2, provided that the average take-rate with this publisher during this billing period is sufficiently high,³⁰⁵ DRSv1 or DRSv2 would instead clear the transaction, charge the winner \$11, collect \$1 and pass on \$10 to the publisher for a fee of approximately 9.1%. DRSv2 further tracks that the advertiser paid \$11 for an impression that ‘costs’ \$12.5 and adds \$1.5 to that advertiser’s debt. DRSv2 also further tracks that the publisher was paid \$10 on a total AdX revenue of \$11, whereas 80% of \$11 would have been \$8.8, and therefore also adds \$1.2 to the publisher’s debt. If these four impressions constitute the same winning advertiser and publisher, then the advertiser has cleared \$1 of debt in the first two impressions and accumulated \$1.5 of debt in the final example, and so will later pay AdX to clear the remaining \$0.5 of debt. The publisher has cleared \$1.2 in the first two

³⁰⁴ Note that the impression might be sold to another exchange or go unsold.

³⁰⁵ If the average take-rate with this publisher is too high, then DRSv1 and DRSv2 both conclude as a regular second-price auction.

impressions and accumulated \$1.2 of debt in the final example, and so is now balanced.³⁰⁶ This example is illustrated in Figure 38 below.

Figure 38: An impression cleared by AdX with the DRSv2 dynamic adjustment of the take rate. Both publisher and advertiser accumulate debt



204. In these examples the advertiser wins the same impressions under DRSv1 and DRSv2 but pays more under DRSv2 (because DRSv2 tracks 'debt' incurred when DRSv1 activates and charges the advertiser to make up for this later).

- As compared to no DRS, the advertiser wins an additional impression, but pays their bid for it under DRSv1 (and under DRSv2, the advertiser not only pays their bid for the impression but further accumulates debt they must pay later!).

³⁰⁶ For this final impression, observe that the advertiser is paying their bid. In particular, if the advertiser had reported a bid of \$12 instead, they would have been charged \$12 (and AdX would have taken \$2, for a fee of approximately 16.66%). If the advertiser had reported a bid of \$10.1, they would have been charged \$10.1 (and AdX would have taken \$0.1, for a fee of approximately 1%). That is, while both DRSv1 and DRSv2 act like a truthful second-price auction for some bids, they are instead a first-price auction in the 'dynamic range' from \$10 to \$12.5, where AdX will charge the winner their bid.

Therefore, the advertiser achieves the same payoff under no DRS and DRSv1, and strictly lower utility under DRSv2 than either.³⁰⁷

- b. As for the publisher, they sell the same impressions to AdX under DRSv1 and DRSv2 but receive \$0.4 less payout under DRSv2 (because DRSv2 tracks ‘debt’ incurred when DRSv1 activates and lowers the take rate and pays less to the publisher later to make up for this). However, different circumstances could have resulted in DRSv2 paying out more than DRSv1 (for example, if the advertiser cleared all their debt in the first two impressions, the publisher would have received the same total payout under DRSv1 and DRSv2). In comparison to no DRS, the publisher receives an additional \$10 from AdX across all impressions under DRSv1, but an unknown amount less from the counterfactual of selling the fourth impression elsewhere. If the publisher would have received more than \$10 for this impression elsewhere, then no DRS yields greater revenues. If they would have received less than \$10, then DRSv1 yields greater revenues. The comparison of DRSv2 to no DRS in this example is nearly identical, except that the publisher loses \$0.4 on the first two impressions (although again, this aspect could have gone differently with different numbers). Finally, observe that because DRSv2 “clears publisher debt”, DRSv2 can “do more DRSv1” (that is, over an extended period of time, there would also be impressions where DRSv1 acts like a clean second-price auction with take rate of 20% and fails to transact the impression through AdX, while DRSv2 instead clears it).

205. Since under DRSv2 the advertisers sometimes pay their bid, DRSv2 corresponds to a dirty second-price auction under some conditions and as a result, DRSv2 is not a truthful auction format. Specifically, whenever DRSv2 charges the winning advertiser their bid and pays the publisher their reserve, it is a dirty second-price auction with a hard reserve of r and a soft reserve of $r/0.8$. This is because, in the naming scheme of the dirty second-price auction definition I

³⁰⁷ Note that the only change between no DRS and DRSv1 is that the advertiser now wins impressions and pays their value. For a sophisticated advertiser who has truly accounted for their budget constraints and knows a value so that they are truly indifferent between paying this value to win the impression or losing, they are truly no better or worse off. For a less sophisticated advertiser who has not fully incorporated budget constraints into their value, they may prefer not to win an impression and pay their value, because this hurts their ‘return on investment (ROI)’ – getting nothing might be preferable to investing \$10 into advertising for a return of exactly \$10.

provide above, when the highest bid b_1 is the unique bid above the hard floor r and below the soft floor $r/0.8$, they win and pay their bid.³⁰⁸

206. DRSv2 goes a step further than being a dirty second-price auction, due to its debt mechanism. In particular, when an advertiser submits a winning bid in the ‘dynamic region’ (*i.e.*, a bid between the publisher-set reserve of r and the effective reserve of $r/0.8$), the advertiser not only pays their bid now resulting in a payoff of 0, but further accumulates debt that must be paid later. That is, *assuming that all debt is eventually cleared*, the auction that advertisers participate in under DRSv2 is as described below, which I’ll call a *debt-aware second-price auction*.

- a. AdX learns its reserve r from the ad server and solicits bids from the ad buying tools. Let b_1 denote the highest received bid and b_2 denote the second highest.³⁰⁹
- b. If the second highest bid is above the effective reserve (*i.e.*, $b_2 \geq r/0.8$), then the clearing price is b_2 .
- c. If the effective reserve is between the highest and second highest bids (*i.e.*, $b_1 \geq r/0.8 > b_2$), then the clearing price is $r/0.8$.
- d. If the highest bid is high enough to clear the reserve, but not the effective reserve (*i.e.*, $r/0.8 > b_1 \geq r$), one of the following happens:
 - i. AdX fails to win the impression. This happens during periods when AdX’s average take-rate during the current billing period with this publisher is much lower than 20%.
 - ii. The clearing price is $r/0.8$, the effective reserve, which strictly exceeds the winning bidder’s bid.³¹⁰
- e. If the highest bid is below the reserve (*i.e.*, $b_1 < r$), no bid is returned and AdX does not win the impression.

³⁰⁸ This is the only conclusion that does not hold for all three payment variants articulated in GOOG-NE-13207241. All of these variants are still not truthful, but do not verbatim fit the ‘dirty second-price auction’ framework. GOOG-NE-13207241. “AdX Dynamic Revshare v2: Launch Doc.” (under the section “Other Pricing Rules”)

³⁰⁹ For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

³¹⁰ This follows because, assuming all debt clears, the advertiser accumulates debt equal to $r/0.8$ minus their payment, which totals $r/0.8$ of debt plus payment.

207. I discuss the implications of participating in a debt-aware second-price auction later in this section, and observe now just that step d.ii. above indeed charges the winning bidder a price that exceeds their bid.³¹¹

C. Truthful Dynamic Revenue Sharing

208. With the third and last iteration tDRS, AdX dynamically adjusted its take rate to sometimes be higher or lower than 20% to win impressions that it would not have if the take rate was kept at 20% on a per-query basis. Under tDRS, Google determined the dynamic take rate it is going to charge before ‘peeking’ at the bids.³¹² In contrast, both DRSv1 and DRSv2 adjusted the take rate after AdX observed the submitted bids.³¹³ Under tDRS, the take rate calculation is done based on the past AdX data. Internal Google documents states that the prediction model that determines the take rate “predicts for a given query whether a specific buyer would bid above the pre-revshare reserve price.”³¹⁴

209. tDRS was fully launched in the second half of 2018.³¹⁵ When AdX migrated to a first-price auction format in 2019, the DRS program was shut off.³¹⁶

[REDACTED]

[REDACTED]

[REDACTED]

³¹¹ And recall again that this follows because we are assuming all debt clears, and so therefore can count payments made to clear debt the moment that debt is accumulated.

³¹² GOOG-AT-MDL-019244499. “Truthful DRS Auction Walkthrough.” (“For each buyer, its reserve price revshare factor will be determined based prediction result before the request is being passed down to RTBs or CAT2 mixer (for Adwords and DBM).”)

³¹³ GOOG-NE-13226622 at -2. “Truthful DRS Design Doc.” (“One known issue with the current DRS is that it makes the auction untruthful as we determine the AdX revshare after seeing buyers' bids and use winner's bid to price itself (first-pricing) when the bid is within the dynamic region.”)

³¹⁴ GOOG-NE-13214748 at -8. “Modeling Design Doc for Truthful DRS.”

³¹⁵ GOOG-TEX-00858434. January 29, 2020. “Dynamic Revenue Share.” (“Update (July 30, 2018): We launched a new DRS model (tDRS).”)

[REDACTED]

[REDACTED]

[REDACTED]

Term	Percentage
Climate change	98%
Global warming	95%
Green energy	92%
Sustainable development	90%
Environmental protection	88%
Green economy	95%

211. Observe that from the advertiser perspective, tDRS is always a regular second-price auction with reserve r^* . As a result, tDRS is truthful, which is different from DRSv1 and DRSv2. This indeed was Google's motivation behind creating tDRS, stating that "the current DRS ... makes the auction untruthful as we determine the AdX revshare after seeing the buyers' bids and use the winner's bid to price itself (first-pricing) when the bid is in the dynamic region. This could incentivize buyers to bid strategically instead of truthfully."³²⁵

212. To illustrate how tDRS works, imagine an impression arrives, and the publisher reports a price floor of \$10. AdX first decides on an alternate reserve of \$9 and a beat-the-floor revenue share of 80%, without seeing any bids. From both the advertisers' and publisher's perspective this executes as a clean second-price auction with a reserve of \$12.5 and a take-rate of 20%.

213. Imagine another impression arrives, and the publisher again reports a price floor of \$10. AdX again decides on an alternate reserve of \$9, but this time a beat-the-floor revenue share of 100%, without seeing any bids. From the advertisers' perspective, this is a second-price auction with reserve \$11.25.³²⁶ From the publisher's perspective, their payout is complicated. The auction will clear whenever b_1 exceeds \$11.25. In this case, if b_2 happens to exceed \$12.5, then the publisher will receive 80% of b_2 , a clean 20% take-rate. If instead b_2 falls below \$12.5, then the publisher will receive \$10, their price floor, and necessarily have a take-rate of < 20% (because both b_2 and the reserve are less than \$12.5, so the winning bidder pays less than \$12.5). If for

³²⁵ GOOG-NE-13226622. "Truthful DRS Design Doc."

³²⁶ This follows as $\$9/0.8 = \11.25 , which exceeds $\$10$.

example b_2 is \$12, then the publisher receives a revenue share of \$10/\$12 for roughly 83.3% > 80%. An 80% revenue share would have given \$9.60, so a debt of \$0.4 is accumulated.³²⁷ If b_2 is less than \$11.25, then the publisher receives a revenue share of \$10/\$11.25 for roughly 88.9% > 80%. An 80% revenue share would have given \$9, so a debt of \$1 is accumulated.³²⁸

214. Imagine a final impression arrives, and the publisher again sets a price floor of \$10. This time, AdX decides on an alternate reserve of \$15 (because AdX believes the impression can fetch bids much higher than \$10), and a beat-the-floor revenue share of 100%.³²⁹ From the advertisers' perspective, this is a second-price auction with reserve \$18.75. From the publisher's perspective, their payout is complicated. The auction will clear whenever b_1 exceeds \$18.75. In this case, if b_2 happens to exceed \$18.75, then the publisher gets 80% of b_2 for a clean 20% take-rate. If b_2 instead falls below \$18.75, then the 80% payout to the publisher would be \$15, but AdX uses this opportunity to clear some of the publisher's 'debt'. Instead of paying the publisher \$15, AdX pays out \$15 less y and clears y 'debt'. y will be at most \$5 (to guarantee that the publisher sees a payment exceeding their price floor of \$10), and also at most \$15-0.8*b₂ (to guarantee that the publisher sees a payment exceeding 80% of b_2).

215. Assuming that all debt is cleared, the publisher receives an 80% revenue share across all auctions AdX clears, and AdX receives payment exactly according to a second-price auction with reserve $r^* = \max\{R/0.8, r/x\}$. In particular, if we count debt when it is accumulated rather than when it is ultimately paid (call this the publisher's debt-aware payout), then the publisher receives 80% revenue share of a second-price auction with reserve $r^* = \max\{R/0.8, r/x\}$ on a per-auction basis. In particular, observe that if $r > R$ and $x > 0.8$, *the publisher's debt-aware payout can be less than their price floor*. In my previous example, when the publisher reports a price floor of \$10, and AdX runs a second-price auction with reserve $r^* = \$11.25$, and $b_2 < \$11.25 < b_1$, the impression clears at a price of \$11.25, and therefore the publisher's debt-aware payout is 80% of \$11.25 = \$9 < \$10.³³⁰

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

³²⁹ When the alternate reserve exceeds the publisher-set reserve, the beat-the-floor revenue share is immaterial.

³³⁰ If tDRS calculated debt according to the formula $r^*(x-0.8)$ instead of $r-\max\{0.8*r/x, 0.8*b_2, R\}$, and only used values of $x = 1.0$ or $x = 0.8$, then tDRS overcalculates debt, and would therefore eventually pay out publishers even less than I compute above. That is, AdX would collect revenue according to a second-price auction with reserve $r^* = \max\{R/0.8, r/x\}$, while giving the publisher at most an 80% debt-aware revenue share, and sometimes strictly less.

D. Impact of DRS Programs on Publishers

216. The impact of DRSv1 on the publishers, as compared to no DRS, has an indeterminate impact on publisher revenue. In a situation where DRSv1 kicks in and helps an impression clear when it would not have otherwise, the publisher is always paid their price floor. This price floor could be greater than, equal to, or less than the revenue the publisher would have gotten had AdX not transacted the impression. For example:

- a. The price floor might be equal to the highest bid received via header bidding. In this case, the publisher would receive exactly the price floor in case AdX does not transact this impression (because the price floor is a bid the publisher has in hand). Therefore, in this case, the publisher is neutral towards DRSv1.
- b. The price floor might be a reserve set on exchanges by the publisher, and it could be that no exchanges cleared this reserve in header bidding. In this case, the outside opportunity if AdX does not transact this impression is zero, because it would otherwise go unsold. In this case, the publisher sees increased revenue from DRSv1.
- c. The price floor might be set via (Enhanced) Dynamic Allocation with other exchanges participating via the waterfall. In this case, the publisher does not know exactly what opportunity would have arisen had AdX not transacted the impression and it instead entered the waterfall. It could be that the publisher would have gotten lucky and seen revenue from the waterfall that exceeds this price floor. In this case, the publisher sees decreased revenue from DRSv1. It could be that the publisher would have gotten unlucky, the impression would not have cleared in the waterfall. In this case, the publisher sees increased revenue from DRSv1. On average, if the publisher sets a reserve strictly above their expected opportunity cost, they would expect to see increased revenue from DRSv1 (although this is not guaranteed to occur impression-by-impression).

217. DRSv2 also has an indeterminate impact on publisher revenue, as compared either to no DRS or DRSv1. DRSv2 has two channels of additional effects, both of which have indeterminate impact: DRSv2 could clear more or fewer transactions in the dynamic region as compared to DRSv1. If advertisers bid truthfully in DRSv2 and DRSv1, then DRSv2 would clear more transactions in the dynamic region, because DRSv2 increases its take-rate on some transactions

so that it can lower its take rate in the dynamic region more frequently. If some advertisers were to realize they are participating in a debt-aware second-price auction and stopped bidding in the dynamic region under DRSv2, then fewer transactions would be cleared in the dynamic region.³³¹ Therefore, if a publisher saw increased revenue from DRSv1 to no DRS, they would also see increased revenue from this aspect of DRSv2 to no DRS. If a publisher saw decreased revenue from DRSv1 to no DRS, they would also see decreased revenue from this aspect of DRSv2 to no DRS. Whether a publisher saw increased revenue from this aspect of DRSv2 to DRSv1 could go either way, depending both on whether they saw increased or decreased revenue from DRSv1 to no DRS, and whether more or fewer impressions transact in the dynamic region.

218. DRSv2 maintains a single debt tracker for each advertiser and each publisher, rather than each (advertiser, publisher) pair.³³² This means that an advertiser can incur debt while purchasing an impression from one publisher (and the publisher incurs 80% of that debt) and clear it while purchasing an impression from another publisher (and this publisher then earns 80% of the cleared debt). This aspect is zero-sum across publishers but causes transfers from some publishers to others. A publisher could see increased or decreased revenues due to this aspect, in comparison to both DRSv1 and no DRS, depending on whether they tend to participate more in debt-building (decreases revenue) or debt-clearing (increases revenue) transactions.

- 1) Publishers would have set different reserve prices to maximize their revenues had Google revealed DRSv1

219. Google concealed material information from publishers by not disclosing the implementation of DRSv1. Since publishers believed that AdX runs a regular second-price with their given reserve price and a static take rate of 20%, a strategic publisher would set a price floor that maximized their revenue under these circumstances. Had they known AdX dynamically adjusted its take rate, publishers would set different price floors.³³³ In general, auction formats are known to be vital for optimizing revenues, as are reserve prices. As a publisher, it is therefore

³³¹ I discuss advertiser incentives under DRSv2 in Section VII.B. In particular, bid-shading in the dynamic region results in the same outcomes as truthful bidding. Outcomes change only if advertisers forego the dynamic region entirely.

³³² See GOOG-NE-13207241 at -45. "AdX Dynamic Revshare v2: Launch Doc." [REDACTED]

[REDACTED]
The example at -46 demonstrates advertiser b_1 incurring debt with publisher p_1 and clearing it with publisher p_2 .

³³³ For example, with a static take rate of 20%, a publisher might naturally set the price floor r such that $r/0.8$ is the revenue-optimal reserve, taking into account the opportunity cost of selling the impression elsewhere. If the publisher knew that the take-rate could go down [REDACTED] on average, a publisher might naturally set the price floor r such that [REDACTED] revenue-optimal reserve (again, taking into account the opportunity cost of selling the impression elsewhere). That is, one natural publisher response to DRSv1 is to increase price floors on AdX.

material to understand what auction AdX is running and how the reserve set on AdX influences the advertiser bids.³³⁴

E. Impact of DRS Programs on Exchanges

220. DRSv1 would lead to an increase in AdX's win rate and an increase in AdX's revenue as compared to no DRS. Since DRSv1 was not disclosed to the advertisers, they would still bid their true values for the impression. The only change that happens with DRSv1 is that sometimes AdX successfully clears an impression when it would have not without DRSv1. Therefore, AdX would win at least every impression that it wins without DRSv1 while paying the same price and might win additional impressions. This leads to an increase its win rate and revenue. This holds whether other exchanges participate via waterfalling or header bidding.³³⁵

221. Internal Google documents show that AdX's win rate and revenue indeed increased as a result of DRSv1. An internal email provides the immediate impact of DRSv1 a week after the launch, stating that it brought an additional (annualized) [REDACTED] in revenue to AdX.³³⁶ Furthermore, the email states "overall match rate for AdX publishers increases by [REDACTED] [REDACTED] when selling to AdX buyers,"³³⁷ demonstrating the increase in the number of transactions cleared by AdX. An internal Google presentation stated that DRSv1 led to [REDACTED] increase in annual recurring revenue.³³⁸

222. DRSv2 would lead to an increase in AdX's win rate and revenue as compared to no DRS. The impact of DRSv2 on AdX win rate and revenue is indeterminate as compared to DRSv1. Because DRSv2, via its debt mechanism, enables AdX to dynamically adjust its take rate more often compared to DRSv1, this aspect would lead to an increase in win rate for AdX compared to DRSv1. Since DRSv1 leads to an increase in AdX win rate compared to no DRS, this means that DRSv2 would lead to a further increase in win rate over no DRS as well. However, since DRSv2

³³⁴ For example, because DRSv1 is material to advertisers' bid decisions, if DRSv1 were to be abruptly disclosed, it could cause advertisers to abruptly start bid shading in AdX, which could abruptly negatively impact a publisher's revenue. As another example, perhaps the publisher is sophisticated and deciding how much to adjust Value CPMs to mitigate AdX's Last Look advantage, which was exacerbated by DRSv1. In such a situation, the publisher might underestimate the impact of AdX's Last Look advantage and mistakenly decide against adjusting Value CPMs further.

³³⁵ These additional impressions might have otherwise been cleared by a non-Google exchange, and therefore that exchange's win rate and revenue would decrease. It is also possible that some of these additional impressions would have otherwise gone unsold, which would not directly impact other exchanges.

³³⁶ GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)."

³³⁷ GOOG-TEX-00777528. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)."

³³⁸ GOOG-NE-06842715 at -18. May 10, 2016. "AdX Auction Optimizations."

is not truthful and aspects of it were disclosed, advertisers may have shaded their bids. I have previously noted that bid-shading itself within the dynamic region itself does not change outcomes under DRSv2, but if some advertisers choose to skip the dynamic region entirely, this could cause fewer impressions to transact in the dynamic region compared to DRSv1. Still, if any advertisers transact in the dynamic region, DRSv2 clears additional impressions beyond no DRS, and therefore DRSv2 leads to increased win rate and revenue for AdX in comparison to no DRS.³³⁹ When compared to DRSv1, AdX's win rate and revenue could increase or decrease, depending on the magnitude of the two effects noted above which push in opposite directions.

223. tDRS would lead to an increase in AdX's revenue and win rate as compared to no DRS. tDRS and no DRS are both truthful auctions, so advertisers would submit the same bids to both. The only distinction between tDRS and DRS is that tDRS gives AdX more flexibility over the effective reserve for its auction (under no DRS, it must be at least $r/0.8$, where r is the publisher's price floor,³⁴⁰ with tDRS, AdX can now set the effective reserve as low as r). AdX would therefore use this flexibility to better optimize its revenue. Moreover, any optimization would come by lowering the effective reserve, which would increase AdX's win rate. Therefore, tDRS would lead to both an increase in AdX's revenue and win rate.³⁴¹

224. DRSv1, DRSv2, and tDRS all naturally exacerbate all conclusions regarding non-AdX exchanges under Dynamic Allocation and Enhanced Dynamic Allocation. In Section IV, I established that exchanges without AdX's advantage under Dynamic Allocation and Enhanced Dynamic Allocation clear fewer impressions and earn less revenue compared to AdX. The gap in win rate and revenue between exchanges with and without AdX's Last Look advantage would be larger when the Last Look advantaged exchange uses some form of DRS in comparison to when all exchanges use fixed take rates. The "Last Look" advantage bestowed upon AdX via Dynamic Allocation and Enhanced Dynamic Allocation gives AdX a reserve price such that if AdX can pay more than this reserve, AdX is guaranteed to win. When other exchanges participate in the waterfall, Dynamic Allocation prevents other exchanges from having an opportunity to return live bids. For Dynamic Allocation with other exchanges in header bidding, this allows AdX to use the

³³⁹ Again, these additional impressions might have otherwise been cleared by a non-Google exchange, and therefore that exchange's win rate and revenue would decrease. It is also possible that some of these additional impressions would have otherwise gone unsold, which would not directly impact other exchanges.

³⁴⁰ Note that a sophisticated publisher could respond to tDRS by increasing their price floors, which would limit AdX's flexibility.

³⁴¹ Again, these additional impressions might have otherwise been cleared by a non-Google exchange, and therefore that exchange's win rate and revenue would decrease. It is also possible that some of these additional impressions would have otherwise gone unsold, which would not directly impact other exchanges.

maximum live bid to inform its reserve. For Enhanced Dynamic Allocation, this allows AdX to use direct deal CPMs as its reserve price. DRSv1, DRSv2 and tDRS give AdX more flexibility to clear this reserve price determined by Dynamic Allocation or Enhanced Dynamic Allocation more often, which further increases Google's win rate and revenue and further decreases the win rate and revenue of other exchanges.³⁴²

F. Impact of DRS Programs on Advertisers

225. DRSv1 was likely neutral to AdX advertisers' payoffs, although may have had some negative impact. First, DRSv1 was concealed, and therefore AdX advertisers would bid their true values into what they believed was a second-price auction. Because advertisers pay their bids when winning an impression in the dynamic region, their payoff is 0 when this occurs (the same as if they did not win). Still, there are two mechanisms by which AdX advertisers may have been negatively impacted. First, if publishers were to regularly re-optimize reserves, their reserves might have naturally increased under DRSv1 (even without fully understanding AdX's auction). If publishers' reserves increase, AdX advertisers pay more for impressions, and therefore suffer negative impact. Second, some advertisers might care about ROI in addition to payoff. Winning an impression that gives payoff 0 harms ROI, and therefore such advertisers would also suffer a negative impact.

1) DRSv2 led to a decrease in AdX advertisers' payoffs

226. DRSv2 was quite negative towards AdX advertisers' payoffs.³⁴³ I previously noted that from AdX advertisers' perspective, DRSv2 is a debt-aware second-price auction. In particular, a debt-aware second-price auction is exactly a second-price auction with reserve $r/0.8$ except if the winning bid is between r and $r/0.8$ it is treated as $r/0.8$ instead. This means three things for how an advertiser with value v should optimally respond: (a) if v is outside the dynamic region (i.e., $v > r/0.8$), the advertiser should optimally bid v , (b) if v is inside the dynamic region, shading v to some other value inside the dynamic region makes no difference,³⁴⁴ (c) if v is inside the dynamic region, the advertiser is best off omitting a bid entirely. Let me now draw a few conclusions.

³⁴² Note that I present this conclusion in Section IV on Dynamic Allocation and Enhanced Dynamic Allocation as well.

³⁴³ At least, those who used third-party ad buying tools, as DV360 and GDN were exempted from DRSv2. GOOG-TEX-00831090 at -1. April 17, 2017. "DRX 2.0 Quality." ("Sellside dynamic revenue share (DRS) is applied to AdX RTB but not Adwords or DBM.")

³⁴⁴ Conditional on this advertiser still winning, bid shading within the dynamic region benefits advertisers only because it makes the advertiser less likely to win in the first place but does not change the outcome if the advertiser still wins. Of course, depending on the extent to which advertisers fully understood the precise mechanics of DRSv2, they may have shaded their bids instead of skipping the dynamic region (depending on how well advertisers fully understood debt repayment, they may have even shaded their bids outside of the dynamic region since bid shading

- a. If all advertisers responded optimally to DRSv2, *no advertiser would bid in the dynamic region, and therefore DRSv2 would be equivalent to no DRS.* That is, exactly the same advertisers would win exactly the same impressions and pay exactly the same amount and the entire DRSv2 program would be obviated. Internal Google documents show that [REDACTED]

[REDACTED]³⁴⁵ This increase in AdX revenue is only possible if transactions are cleared in the dynamic region, and every transaction cleared in the dynamic region necessarily involves an advertiser paying more than their value for an impression, and therefore suffering decreased payoff in comparison to no DRS. I again wish to emphasize that any increase in AdX revenue *must* come at the cost of some advertisers having lower payoff under DRSv2 compared to no DRS, *and those advertisers behaving sub-optimally under DRSv2.*

- b. No matter how an advertiser responds to DRSv2, they cannot possibly be better off than no DRS. In particular, the best they can do is to bid truthfully outside the dynamic region and avoid the dynamic region entirely. Doing so will give them the exact same outcomes as under no DRS, and any other behavior causes a decrease in their payoff as compared to no DRS.
- c. Bid-shading in the dynamic region does not improve an advertiser's outcomes over truthful bidding,³⁴⁶ they must skip the dynamic region entirely in order to better respond to DRSv2 over truthful bidding. This fact only follows from the precise mechanism that DRSv2 uses to track and clear debt and cannot be deduced merely from the fact that DRSv2 dynamically increases and decreases AdX's take rate.

outside the dynamic region does cause advertisers to pay less when they win, but this is only because they are clearing less debt). Moreover, if advertisers did not know for which auctions DRSv2 was active, advertisers who considered bid shading may have further shaded their bids even on AdX auctions that concluded as truthful second-price auctions without DRSv2.

³⁴⁵ GOOG-NE-13234466 at -67. "Overall Pub Yield With DRS(v2)."

³⁴⁶ Again, bid shading could only help by reducing the risk that the advertiser wins, or causing the advertiser to transact in the dynamic region less often. Conditioned on winning the same set of impressions, bid shading has no impact.

- 2) Advertisers would have submitted different bids to maximize their payoffs had Google revealed DRSv1

227. Google misled advertisers by not revealing DRSv1, and hence led them to believe the AdX auction was a regular second-price auction, which would cause them to engage in suboptimal behavior. When advertisers believe they are participating in a regular second-price auction, they would bid their true value for the impression, because it is a truthful auction. However, DRSv1 is not truthful, as established before. Therefore, concealing DRSv1 caused advertisers to bid their true value in a non-truthful auction, whereas advertisers would get higher a higher gain by shading their bids.

228. By not revealing DRSv1 to the advertisers, Google made material gains. This is because if advertisers were to shade their bids, which is the natural bidding behavior in a non-truthful auction like DRSv1, this would lead to less revenue for both AdX and publishers.³⁴⁷ However, advertisers likely did not shade their bids, since Google never publicly revealed DRSv1.

G. Some aspects of DRS are exceptionally misleading

229. To conclude the section, I want to briefly note a few aspects of DRSv2 that I find exceptionally misleading to advertisers, and an aspect of DRSv2 and tDRS that I find misleading to publishers. Much of my analysis below concerns the concept of ‘debt’ to mislead both advertisers and publishers regarding how much they are paying or paid out.

230. First, I want to repeat that my previous analysis establishes that when advertisers behave optimally in DRSv2, no transactions should ever occur in the dynamic region. Instead, Google claims that enough transactions occurred in the dynamic region to account for [REDACTED]
[REDACTED].³⁴⁸ This increase necessarily comes at the expense of advertisers ultimately *paying more than their value for an impression*.

231. Next, I want to highlight aspects of Google’s description³⁴⁹ regarding DRS that I find misleading, due to the concept of debt.

³⁴⁷ But this does not mean DRSv1 as a whole, after accounting for bid shading, would necessarily lose revenue. It merely means that when comparing “DRSv1 when advertisers shade their bids” yields higher payoff for advertisers and less revenue for AdX and the publishers than “DRSv1 when advertisers bid their true values,” which motivates concealing DRSv1.

³⁴⁸ GOOG-NE-13234466 at -67. “Overall Pub Yield With DRS(v2).”

³⁴⁹ GOOG-NE-04934281 at -84. July 30, 2018. “Dynamic Revenue Share.”

- a. For DRSv2, Google states: “Buyers are never charged more than their bid,” I find this claim exceptionally misleading to advertisers. If by ‘charged’, one means ‘charged on this impression as immediate payment, ignoring any debt that will be paid later’, then the sentence is technically true. But if by ‘charged’, one means ‘charged on this impression either as immediate payment or as debt to be collected later’, then the sentence is false. Indeed, any winner in the dynamic region is ultimately charged more than their bid after accounting for both immediate payment and debt to be collected later.
- b. For DRSv2, Google states: “sellers are always paid at least their reserve.” I find this claim misleading to publishers. If by ‘paid’, one means ‘paid on this impression as immediate payment, ignoring any debt generated that must be paid back later’, then the sentence is technically true. But if by ‘paid’, one means ‘paid on this impression as immediate payment, after subtracting any debt assigned that will be collected later’, then the sentence is sometimes false. On impressions transacted in the dynamic region, publishers are indeed paid their reserve as immediate payment, but are also assigned non-negative debt. This debt may wind up being cancelled (if the winning advertiser clears their own debt while later transacting with this publisher), owed (if the winning advertiser clears their own debt elsewhere), or yielding extra payout (if some other advertiser later clears debt incurred elsewhere with this publisher). Even after accounting for debt, it is plausible to say that “sellers as a whole are paid at least the reserve set on any cleared impression.” But it is inaccurate to say “sellers are *always* paid at least their reserve,” because some sellers are not ultimately paid their reserve.³⁵⁰
- c. If the concept of debt was not clearly disclosed, the general description of DRS as per-query revenue share optimization is insufficient for advertisers to draw conclusions at the level I have drawn in my report. Moreover, even for advertisers who are already optimizing bids at a per-impression level, the concept of debt significantly obscures feedback. Indeed, a typical optimizer might ask questions of the form “if I change my bid on this impression, what change does that cause in my payoff from this impression?” The concept of debt now means that changing a

³⁵⁰ Sellers would reasonably care about whether they are paid their reserve or not, as this reserve constitutes the minimum amount they’ve decided to accept in order to forego the opportunity cost of selling the impression elsewhere.

bidder's bid on an impression *might cause them to pay more on another impression* (when that debt is cleared). In order for advertisers to have sufficient information regarding DRSv2 in order to avoid paying more than their value for an impression, Google would have needed to disclose a somewhat precise description of the debt concept. I do not know whether Google indeed made such a disclosure, nor how it was made, but in my opinion such information is vital to advertisers, even if they were already aware in a general sense that DRSv2 optimizes revenue shares on a per-query basis.

- d. In a Google communications document,³⁵¹ that seems to have been active after tDRS launched,³⁵² Google states "Before or after the July change, we still do NOT pay publishers below publisher's floor..." I find this claim quite misleading to publishers concerning tDRS. If by 'pay', one means 'pay as immediate payment, ignoring any debt generated that must be paid back later', then the sentence is technically true. But if by 'pay', one means 'pay on this impression as immediate payment, after subtracting any debt assigned that will be collected later', then the sentence is false. Indeed, after accounting for debt assigned, I previously noted that the debt-aware payment to the publisher is sometimes less than their price floor. Even after accounting for debt, it is fair to say that "any publisher can guarantee a debt-aware payment at least as high as a desired price floor r by setting AdX's price floor to 1.25^*r ." But I find it misleading to blanketly assert that the publisher is not paid below their floor without further nuance.³⁵³

VIII. CONDUCT ANALYSIS: PROJECT BERNANKE

232. In this section, I analyze the conduct referred to as Project Bernanke, which was later expanded to Project Global Bernanke, and Project First Price Bernanke (also called "The Alchemist"). I also draw conclusions regarding their effects. I show that that Projects Bernanke and Global Bernanke did not affect GDN advertisers' payoffs and could increase some publishers' revenues while decreasing others, however they led to a lower win rate for non-GDN ad buying tools and advertisers that use those ad buying tools. Furthermore, Project Bernanke and its variants led to an increased win rate for GDN buyers (and in the case of Projects Bernanke and

³⁵¹ GOOG-TEX-00858434 at -40. January 29, 2020. "Dynamic Revenue Share."

³⁵² The document seems to be updated through at least December 2019, which is after the tDRS launch.

³⁵³ Sellers would reasonably care about whether they are paid their floor or not, as this floor constitutes the minimum amount they have decided to accept in order to forego the opportunity cost of selling the impression elsewhere.

Global Bernanke, without improving GDN buyers' utilities), which leads to an increased win rate and revenue for GDN (again in the case of Projects Bernanke and Global Bernanke, without assisting GDN advertisers at all).

233. Project Bernanke and all its variants can be understood as simultaneously facilitating the effects of collusion among GDN advertisers, without their knowledge, and overbidding in auctions. I explain this view in detail in this section. The starting point for Project Bernanke is Google's observation that [REDACTED]

[REDACTED].³⁵⁴ Because GDN is 'second-pricing itself', GDN would benefit by lowering the second-highest bid it sends in order to lower the payment GDN must make to win the impression.³⁵⁵ In isolation, this would be a pure transfer of funds from exchange/publisher to GDN, but would technically result in a high take rate by GDN towards its advertisers, compared to what is contracted. This aspect of Project Bernanke is akin to facilitating collusion among GDN advertisers (and in this case, without their knowledge).³⁵⁶ The second half of Project Bernanke uses the savings from the first half and spends it to subsidize overbidding.³⁵⁷ That is, the second half of Project Bernanke boosts the bids of its advertisers before sending them to AdX, but uses the funds from the first half to cover any payment made above the advertiser's true bid. In all Project Bernanke variants, the two halves balance out to (a) generate increased revenue and increased win rate for GDN, (b) balance GDN's take rate at the intended 14%,³⁵⁸ (c) have an indeterminate effect on publisher revenue (the first half decreases publisher revenue while the second half increases it), (d) in the case of Project Bernanke and Project Global Bernanke, not improve GDN advertisers' payoffs at all. Below I describe Project Bernanke in greater detail, and Appendix H contains additional discussion on overbidding and collusion in first- and second-price auctions.

³⁵⁵ GOOG-AT-MDL-001412616 at -19.

³⁵⁶ By using the word ‘collusion’, I do not mean to imply that it is ‘wrong’ from a pure auction theory perspective for a group of bidders to get together and jointly strategize on how to collectively bid, nor for an ad buying tool to facilitate this. My understanding is that other ad buying tools may have dropped their second-highest bid entirely. GOOG-NE-13200831 at -1. “The case for encouraging buyers to declare two bids.” (“Currently, the only [AdX] buyer who is employing this strategy [sending two bids] is GDN.”)

Still, GDN is indeed facilitating collusion by implementing a joint strategy for its advertiser pool together, rather than processing each advertiser's bid in isolation.

³⁵⁷ GOOG-AT-MDL-001412616 at -20. "Project Bernanke and margins story." ("What if overbid? We could bid too much. But we have to subsidize it. One is good for us [GDN] and bad for publishers. Other is bad for us [GDN] and good for publishers.")

³⁵⁸ Some documents I have reviewed state that the GDN take rate is 14% and others state that it was 15%. I use 14% throughout the text, except in the cases where I cite a specific document that states 15%. This difference in the take rate does not have any effect on my conclusions throughout this section.

A. Project Bernanke

234. Prior to any modifications, GDN ran an internal auction (called the “CAT2” auction)³⁵⁹ with only GDN advertisers and submitted the top two bids from that auction to AdX.^{360, 361} Under **Project Bernanke**, between 2013 and 2015,³⁶² GDN manipulated advertisers’ bids before sending them to AdX in the following manner:³⁶³

a.

Category	Value
1	~95
2	~98
3	~90
4	~98
5	~95
6	~98
7	~95
8	~98
9	~95
10	~98
11	~95
12	~98

³⁵⁹ Also called the “CAT2 auction.” GOOG-NE-11753797 at -37. February 11, 2019. “DVAA Quality, Formats, O&O - Q1 2019 All Hands.”

³⁶⁰ GOOG-NE-06839089 at -94. "Project Bernanke." ("GDN submits two bids into AdX auction...")

³⁶¹ During the entire lifetime of Project Bernanke, AdX conducted second-price auctions. This subsection covers the original Project Bernanke, and later subsections cover its variants.

³⁶² Project Bernanke was launched in 2013. It was in place until Project Global Bernanke was launched in 2015. See GOOG-DOJ-28385887 at -93, 94. August 17, 2015. "Beyond Bernanke." ("Bernanke (late 2013)... Global Bernanke

³⁶⁵ To be clear, I am opining that this aspect of Project Bernanke, and collusion among GDN bidders, helps GDN at the expense of publishers. This is not an opinion on Project Bernanke as a whole.

³⁶⁵ To be clear, I am opining that this aspect of Project Bernanke, and collusion among GDN bidders, helps GDN at the expense of publishers. This is not an opinion on Project Bernanke as a whole.

³⁶⁷ To be clear, I am opining that this aspect of Project Bernanke, and overbidding in a second-price auction, helps publishers at the expense of GDN advertisers – this is not an opinion on Project Bernanke as a whole.

ii. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]


235. Project Bernanke computes the pair of adjustment parameters (α, β) using historical GDN data³⁶⁸ and auction simulations³⁶⁹ on a per-publisher basis in a manner that maintains an average take rate of 14% for each billing period. In order to maintain an average take-rate of 14%, Google created for each publisher a “Bernanke pool” that is added to whenever GDN’s take rate exceeds 14% and consumed from whenever GDN’s take rate falls below 14%.

236. An excerpt in Figure 39 from an internal Google document below shows how GDN formulated Project Bernanke to maximize GDN revenue while maintaining a take rate of 15% per publisher.

³⁶⁸ GOOG-NE-13468541 at -42. “Bernanke experiment analysis.” (“It is important to note that in this entire process, we only use information about the GDN bid and the GDN price paid on queries won by GDN.”)

³⁶⁹ GOOG-NE-13468541 at -42. “Bernanke experiment analysis.” (“The optimal combination of first bid increase and second bid decrease for each publisher is estimated using AdX auction simulations...”)

Figure 39: An excerpt from a Google slide deck stating that α and β were chosen to maximize GDN revenue while maintaining a take rate of 15%³⁷⁰



237. Project Bernanke results in an increased win rate and revenue for GDN. I explain why this is the case below as well. Because GDN submits a higher top bid to AdX under Project Bernanke, this means that GDN clears more impressions and therefore has a higher win rate. Moreover, Project Bernanke collects a weakly higher payment from its top advertiser on every single impression. In particular, if GDN would have cleared the impression without Project Bernanke, it collects exactly the same payment. If GDN clears the impression only due to Project Bernanke, it collects non-zero payment (as opposed to zero without Project Bernanke). Because Project Bernanke balances GDN's revenue as 14% of its received payment, GDN also sees increased revenues.



238. The excerpts in Figure 40 from an internal Google document³⁷¹ present examples of how Project Bernanke works. In both slides, GDN's top bid into AdX is raised to \$1.20, even though the top GDN value is \$1.00. This can be interpreted as overbidding, which withdraws money from the Bernanke pool. In both slides, GDN's second bid into AdX is decreased to \$0.48 instead of \$0.83 (if GDN submitted the true value of the second-highest bidder less a 14% take-rate on

³⁷⁰ GOOG-NE-11753797 at -37. February 11, 2019. "DVAA Quality, Formats, O&O - Q1 2019 All Hands." (at the time of this document in 2019, GDN take rate was 15%).

³⁷¹ GOOG-NE-06839089 at -98. October 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

\$0.96, this would be \$0.83). This can be interpreted as collusion. Lowering the second-highest bid, after learning that there is a higher GDN bid, helps the highest GDN bidder. In “Project Bernanke: Example #1”, overbidding is necessary to win, as without it the top GDN bidder would not have cleared the publisher’s reserve of \$1.10. In this case, GDN must pay AdX \$1.10, but can only charge its top advertiser at most \$1.00, and therefore this transaction consumes \$0.10 from the ‘Bernanke pool’ to subsidize overbidding. In “Project Bernanke: Example #2”, GDN second-prices itself, and the collusion has a positive impact on GDN. Instead of paying AdX \$0.83, GDN only pays AdX \$0.48 and can pocket the difference into its ‘Bernanke pool’ to subsidize more overbidding. In the long run, Google structured Project Bernanke so that these additions and withdrawals cancel one another out, allowing GDN to win auctions it otherwise would not have won. Appendix H contains numerical examples of these dynamics.

Figure 40: Excerpts from an internal Google slide deck show that Project Bernanke dynamically adjusts the highest and second highest GDN bids before sending them to the AdX auction³⁷²



B. Project Global Bernanke

239. Project Bernanke was later expanded to **Project Global Bernanke** in 2015,³⁷³ which instead targeted an average take-rate of 14% across all publishers [REDACTED]

[REDACTED]
[REDACTED] 374

240. Project Global Bernanke maintained a single pool for all publishers across AdX instead of individual per-publisher pools. Project Global Bernanke aimed to keep this single pool roughly empty at the end of each billing period, implying an average take rate of 14% across all of AdX, subject to the aforementioned additional constraints such as floors on the revenue of individual publishers and on the average take rate for individual publishers.

241. Project Global Bernanke also certainly increased GDN's win rate and revenues, for exactly the same reasons as Project Bernanke. I elaborate on why this is the case below.

C. Impact of Projects Bernanke and Global Bernanke on Publishers

242. Under Project Bernanke, individual publisher revenues may increase, may decrease, or may stay the same. The impact on a publisher depends predominantly on whether Project Bernanke results more often in 'poaching' an impression from other ad buying tools or in clearing impressions that otherwise would have gone unsold because no bids surpassed the reservation price.

243. More specifically, with per-publisher pools, when Project Bernanke's overbidding causes GDN to win an impression it otherwise would not have, the impact of these impressions on a publisher depends on the comparison of (a) the payment AdX would have received without Project Bernanke, and (b) the highest pre-Bernanke GDN bid for this impression. If (a) exceeds (b) on average (*i.e.*, if Project Bernanke largely results in GDN 'poaching' impressions that would otherwise have been won by other ad buying tools), then the publisher faces a revenue loss. If (b) exceeds (a) on average (*i.e.*, if Project Bernanke largely results in GDN winning impressions that otherwise would have gone unsold), then the publisher revenue increases. This is because each publisher loses revenue when the Bernanke pool grows (and by exactly the amount the pool grows), and gains revenue when the Bernanke pool shrinks (and by exactly the amount the pool shrinks), plus the difference between (b) and (a).³⁷⁵ Because Project Bernanke constrains each publisher to have a neutral pool on average, the cumulative average increase/decrease to the

³⁷² GOOG-NE-06839089 at -98, 99. October 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

³⁷³ Also called Project Bell v1. See GOOG-AT-MDL-006218257 at -63. December 16, 2022. "Case AT.40670 - Google - Adtech and Data-related practices." ("Project Bell was also known as Global Bernanke.")

³⁷⁴ GOOG-DOJ-AT-02471194 at -4. July 26, 2015. [REDACTED]

³⁷⁵ See Tables 1, 2 and 3 in Appendix H.5.

publisher's revenue from these impressions comes exactly via the difference between (b) and (a).³⁷⁶ Therefore, the change in revenue earned by publishers through impressions that GDN wins through AdX is determined exactly via the difference between (b) and (a).

1) Project Global Bernanke led to revenue reduction in some publishers

244. Under Project Global Bernanke, individual publisher revenues may increase, may decrease, or may stay the same. The reasoning is the same as the paragraph above, plus one additional complexity due to the global pool. Specifically, Project Global Bernanke creates a further complication with its single global pool under which some publishers may incur losses simply because they contribute more to the Global Bernanke pool than the payouts they receive from it. Stated differently, some publishers may see decreased revenues by Project Global Bernanke because their impressions may have contributed to the Bernanke pool on average, whereas other publishers may have consumed from the pool on average. As a result, it is still the case that publishers for whom (a) is larger than (b) tend to see decreased revenues under Project Global Bernanke, and those for whom (b) is larger than (a) tend to see increased revenues. As compared to the original Project Bernanke, under Project Global Bernanke it is also the case that publishers who tend to contribute to the Global Bernanke pool (*i.e.*, AdX often provides both the highest and second highest bid for such publishers' impressions) more than they consume (*i.e.*, AdX often is not the highest bid for such publishers' impressions) may also see decreased revenues under Project Global Bernanke compared to Project Bernanke, whereas those publishers that consume more than they contribute tend to see increased revenues from Project Global Bernanke as compared to Project Bernanke.

245. In sum, the above arguments imply that:

- a. When comparing the revenues a publisher sees under Project Bernanke versus no Bernanke, this predominantly comes down to comparing the average across all impressions cleared due to Project Bernanke of (a) payment AdX would have received without Project Bernanke and (b) the highest pre-Bernanke GDN bid.

³⁷⁶ In addition, Projects Bernanke and Global Bernanke can also cause a non-GDN advertiser to pay more for an impression that they still win under Projects Bernanke and Global Bernanke. For example, imagine that a non-GDN advertiser bids \$10 on an impression, and the second-highest bid is \$4 from GDN (without Project Bernanke). With Project Bernanke, the second-highest bid might be increased to \$8. This is not high enough to claim the impression, but causes the non-GDN advertiser to pay more, and the publisher to claim an extra \$4 in revenue.

Publisher revenues predominantly decrease if (a) is higher than (b), and predominantly increase if (b) is higher than (a).³⁷⁷

- b. When comparing the revenues a publisher sees under Project Global Bernanke versus Project Bernanke, this is predominantly determined by comparing whether the publisher contributes more than they consume from the Global Bernanke pool.³⁷⁸ Publisher revenues predominantly decrease if they contribute more than they consume, and predominantly increase if they consume more than they contribute.
- c. When comparing the revenues a publisher sees under Project Global Bernanke versus no Bernanke, a publisher for whom (a) is greater than (b) and who contributes more than they consume certainly sees predominantly decreased revenues, and a publisher for whom (b) is greater than (a) and who consumes more than they contribute certainly sees predominantly increased revenues. A publisher for whom (a) is greater than (b) but who consumes more than they contribute, or for whom (b) is greater than (a) but contributes more than they consume, would require a quantitative comparison to ultimately determine their change in revenue from no Bernanke to Project Global Bernanke.

246. Google internal documents show these disparate effects of Project Global Bernanke on publishers as compared to Project Bernanke. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

³⁷⁷ As previously noted, another relevant factor is the increased revenue due to Project Bernanke increasing the payment of non-GDN winners. Put another way, the revenue paid by GDN to AdX increases or decreases entirely based on a comparison of (a) to (b), but the payout of non-GDN advertisers to AdX could increase.

³⁷⁸ An additional effect is that Project Global Bernanke formulates a global optimization rather than per-publisher optimizations, so the set of impressions impacted by Project Global Bernanke will change, and the precise multipliers will also change, as compared to Project Bernanke. For example, due to the flexibility of global versus per-publisher optimization, more impressions were likely impacted by Project Global Bernanke than Project Bernanke. I say 'predominantly' to indicate that this is the novel form of impact when comparing Project Bernanke to Project Global Bernanke, but to not preclude other forms of impact.

³⁷⁹ GOOG-DOJ-AT-02471194. July 26, 2015. "Global Bernanke."

- 2) Projects Bernanke and Global Bernanke can lead to a reduction in ad quality as well as revenue per mille for publishers

247. If GDN advertisers tend to display lower quality ads, Projects Bernanke and Global Bernanke would lead to lower quality ads displayed for publishers. Because GDN wins more impressions under Projects Bernanke and Global Bernanke, if GDN advertisers tend to display lower quality ads than non-GDN advertisers, then this program would cause publishers to receive lower quality ads because it causes GDN advertisers to win impressions more frequently.

248. Even publishers who saw a revenue increase under Projects Bernanke and Global Bernanke can see their revenue per mille (“RPM”) decreased. As shown by an internal Google presentation,³⁸⁰ Project Bernanke led to a decrease in RPM across publishers. This decrease would be because a publisher can benefit from Projects Bernanke and Global Bernanke on a per-auction basis only when an impression is sold that would otherwise have gone unsold because no bids exceeded the reservation price. Hence, a publisher who is concerned about RPM might still prefer the outcomes under no Bernanke as compared to Projects Bernanke and Global Bernanke, even if their overall revenue does not decrease.

- 3) Publishers would have raised their reserve prices to maximize their revenue had they known about Projects Bernanke and Global Bernanke

249. Google concealed vital information from publishers by concealing Projects Bernanke and Global Bernanke. In particular, there are two tests a publisher could do to determine how Projects Bernanke and Global Bernanke predominantly impact their revenues compared to no Bernanke.

- a. As noted earlier, publisher revenue might increase, decrease, or be neutral under Projects Bernanke and Global Bernanke. Under Project Bernanke, a publisher benefits on impressions purchased by GDN only when Project Bernanke largely results in GDN winning impressions that otherwise would have gone unsold rather than in GDN winning inventory where a non-GDN buyer previously set the clearing price. A publisher with a low rate of unsold inventory, or a high volume of competitive non-GDN bids, would likely see decreased revenues under Project Bernanke and prefer their outcomes without Project Bernanke.
- b. Additionally, under Project Global Bernanke, publishers who tend to have GDN provide both the highest and second highest bid more often tend to contribute to

³⁸⁰ GOOG-DOJ-28386151 at -61. December 10, 2013. “Project Bernanke - Quantitative Easing on the AdExchange.”

the Global Bernanke pool rather than consuming from it.³⁸¹ Such publishers would prefer their outcomes without Project Global Bernanke.

250. Moreover, all publishers likely would have changed their behavior if they knew about Projects Bernanke and Global Bernanke by raising their reserve prices.³⁸² For a publisher to see predominantly increased revenues from Project Bernanke, they must tend to interact with it primarily through GDN purchasing unsold inventory rather than GDN stealing impressions from another ad buying tools, and for Project Global Bernanke, they must further interact primarily through GDN purchasing unsold inventory rather than GDN submitting the top two bids. A publisher who raises their reserves causes (a) more inventory to go unsold, increasing the positive interactions with Project Bernanke, (b) non-Google ad buying tools to win less often (because they may fail to clear the reserve), and (c) GDN to be less likely to submit top two bids (because the second price is less likely to clear the reserve). Hence, raising reserves would cause a publisher to primarily have revenue-boosting rather than revenue-draining interactions with Project Bernanke. Under Project Global Bernanke, raising reserves would further cause a publisher to primarily have pool-consuming interactions (which increase publisher revenue) rather than pool-contributing interactions (which decrease publisher revenue) with Project Global Bernanke.

D. Impact of Projects Bernanke and Global Bernanke on Ad Buying Tools

- 1) Under Projects Bernanke and Global Bernanke, GDN increased its revenue at the expense of non-Google ad buying tools

251. Under Project Bernanke and Global Bernanke, GDN increases its revenues and win rate. GDN maintains an average 14% take rate of all revenue paid by its advertisers in AdX auctions. Under Projects Bernanke and Global Bernanke, its advertisers make weakly higher payments on every single instance (which therefore increases GDN's win rate) because:

- a. If GDN would have won without Projects Bernanke or Global Bernanke, the winning GDN advertiser pays the same. The only difference under Project Bernanke is that some portion of this payment now goes to the Bernanke pool.

³⁸¹ Whereas publishers who tend to sell impressions that GDN would not win without Project Global Bernanke (either because a non-GDN advertiser wins, or because the impression would go unsold) would consume more than they contribute.

³⁸² Importantly, Reserve Price Optimization, which I analyze in Section IX, does not apply to GDN. GOOG-TEX-00831090 at -1. April 17, 2017. "DRX 2.0 Quality." (see table.)

b. If GDN would not have won without Projects Bernanke or Global Bernanke, the winning GDN advertiser pays their bid which is certainly larger than \$0 they would have paid if there were no Project Bernanke or Global Bernanke.

252. Google internal documents provides measurements of GDN benefits created by Projects Bernanke and Global Bernanke. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

253. Under Projects Bernanke and Global Bernanke, the win rate of non-GDN advertisers and non-GDN ad buying tools on AdX would decrease. This follows immediately because GDN submits a weakly higher bid for every AdX impression under Projects Bernanke and Global Bernanke as compared to no Bernanke. In addition, non-GDN advertisers would pay more for any impressions they still win under Projects Bernanke and Global Bernanke, for the same reason, but their total aggregate payment could still decrease due to winning fewer impressions.³⁸⁶

254. Google internal documents confirm this analysis as well. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

³⁸³ GOOG-NE-03872763 at -81. "Discussion on improving AdX and AdSense backfill."

³⁸⁴ GOOG-DOJ-28385887 at -95. August 17, 2015. "Beyond Bernanke."

³⁸⁵ GOOG-DOJ-28386151 at -60. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

³⁸⁶ GDN impacts non-GDN advertisers and non-ad buying tools only by determining their minimum bid to win. Only GDN's highest submitted bid impacts others' minimum bid to win, and others' minimum bid to win increases with GDN's highest submitted bid. Once non-GDN advertisers have a higher minimum bid to win, they're immediately less likely to win, and also will pay more if they still win.

³⁸⁷ GOOG-DOJ-AT-02513569 at -73. "gTrade Team Background."

³⁸⁸ GOOG-DOJ-28386151 at -69. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

E. Impact of Projects Bernanke and Global Bernanke on Advertisers

- 1) Projects Bernanke and Global Bernanke did not benefit GDN advertisers, but decreased win rates for advertisers using non-Google ad buying tools

255. Under Projects Bernanke and Global Bernanke compared to no Bernanke, I would expect GDN advertisers to achieve identical payoff, defined as the difference between their value for impressions won and payments made.³⁸⁹ In particular, the only change that GDN advertisers see under Project Bernanke and Global Bernanke is that they now sometimes win impressions and pay their value (for a net payoff of \$0) whereas they would have lost without Project Bernanke (which yields a payoff of \$0).³⁹⁰ More specifically,

- a. In auctions where Projects Bernanke and Global Bernanke decrease the second highest GDN bid to replenish the Bernanke pool, the advertiser wins and pays the second highest bid in full (even though the amount GDN pays to AdX decreases). Had there been no manipulation, the advertiser would have still won and paid the second highest bid again. As a result, these auctions have no impact on advertisers.
- b. In auctions where GDN wins only due to Projects Bernanke or Global Bernanke increasing its highest bid, the GDN advertiser wins and pays their bid, leading to a payoff of \$0.³⁹¹ Had there been no manipulation, they would have lost the impression, either to some non-GDN advertiser or because their bid was below the reserve, and paid nothing. This would lead to a payoff of \$0 as well. As a result, these auctions do not impact GDN advertisers' payoffs either.

Hence, overall, Projects Bernanke and Global Bernanke were neutral towards GDN advertisers' payoffs.³⁹²

³⁸⁹ This subsection assumes that advertisers bid their true value, which is a dominant strategy in a second-price auction with reserve (see Section II). Because Projects Bernanke and Global Bernanke were never revealed to advertisers, this is reasonable behavior to expect. The conclusions I draw still hold approximately if instead some advertisers slightly shaded their bids in response to Bernanke or Global Bernanke.

³⁹⁰ If an advertiser is interested in Return on Investment (ROI) in addition to payoff, that advertiser would prefer no Bernanke to Project Bernanke or Global Bernanke. Specifically, winning an impression valued at v for price v is neutral towards payoff, but harms average ROI (because the ROI on this impression is as low as possible while paying at most the value).

³⁹¹ Again, this assumes that GDN advertisers report their true value to GDN, which is the dominant strategy in a second-price auction (and AdX was described as a second-price auction).

³⁹² There are two caveats to this. First, advertisers who considered shading the bid they entered to GDN, even though AdX was described as a (truthful) second-price auction, may have seen increased payoff. Second, advertisers who value ROI in addition to payoff could still have preferred outcomes without Projects Bernanke or Global Bernanke, even with a neutral effect on payoff.

256. Under Projects Bernanke and Global Bernanke, the win rate of non-GDN advertisers on AdX would decrease. This is because the GDN advertisers are still winning every impression that they would have won without Projects Bernanke and Global Bernanke, but they are also winning additional impressions. Some of these impressions previously would have been won by non-GDN advertisers, so these advertisers face a lower win rate.³⁹³

257. Google internal documentation shows that non-GDN advertisers saw a decline in their win rates. [REDACTED]

[REDACTED] This decrease in spending can only be explained by a decrease in their win rates.³⁹⁶

- 2) Advertisers would have shaded their bids to maximize their payoff had they known about Projects Bernanke and Global Bernanke

258. Google concealed vital information from advertisers by concealing Projects Bernanke and Global Bernanke. Provided that neither Project Bernanke nor Project Global Bernanke were disclosed to advertisers, they would naturally believe they were still participating in a truthful second-price auction and bid their true value as a result. If advertisers knew they were participating in a non-truthful auction, they would have instead considered shading their bids. Knowing the auction format is vital information to advertisers aiming to optimize their payoff. In particular, Projects Bernanke and Global Bernanke are dirty second-price auctions.³⁹⁷ Specifically, if c denotes the minimum bid to win for GDN on AdX, then from the perspective of a GDN advertiser, Projects Bernanke and Global Bernanke are both dirty second-price auctions with soft floor c and hard floor c/α . That is, as long as the highest GDN bidder exceeds c/α , they will win, because their bid will be increased by α to exceed c . If their bid further exceeds c , the auction turns into a regular second-price auction. If their bid falls between c/α and c , they will pay their bid and be subsidized by Bernanke pool for the remaining amount. Therefore, this is a dirty second-price auction.

³⁹³ Some of these impressions could have been previously unsold.

³⁹⁴ GOOG-DOJ-28385887 at -95. August 17, 2015. "Beyond Bernanke."

³⁹⁵ GOOG-DOJ-28386151 at -67. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."

³⁹⁶ This is because, as noted previously, the price a non-GDN advertiser pays if they still win under Projects Bernanke and Global Bernanke can only increase from what they would pay with no Bernanke.

³⁹⁷ See Section VII for more information on dirty second-price auctions.

259. Moreover, it is possible that some advertisers preferred outcomes under no Bernanke as compared to Projects Bernanke or Global Bernanke. Advertisers might be concerned about ROI in addition to payoff. Projects Bernanke and Global Bernanke result in lower ROI for GDN advertisers, because GDN advertisers spend more but without generating positive returns (by purchasing impressions at a price exactly equal to their willingness to pay). Therefore, an advertiser who prefers a higher ROI option among two with equal payoffs would prefer no Bernanke to Projects Bernanke or Global Bernanke.

260. In addition, advertisers may choose to use GDN simply to avoid the negative impacts caused by Projects Bernanke and Global Bernanke to non-GDN advertisers, and not because of any benefits they receive from Project Bernanke or Global Bernanke. In particular, I have previously described that Projects Bernanke and Global Bernanke are neutral towards GDN advertisers' payoffs, and only harm ROI. On the other hand, I have also previously described that Projects Bernanke and Global Bernanke decrease non-GDN advertisers' payoffs (both by decreasing their win rate, and by increasing their payment per-impression). Therefore, had they known about the conducts, an advertiser might prefer GDN to non-GDN simply to avoid the negative impacts caused to non-GDN advertisers by Projects Bernanke and Global Bernanke.

261. Under Projects Bernanke and Global Bernanke, market efficiency may increase, may decrease, or may be stable. Overall efficiency is maximized when the impression is awarded to the highest value bidder. Projects Bernanke and Global Bernanke affects this in two ways. First, if the impression would have been sold to a non-GDN advertiser but is now awarded to a lower value GDN advertiser, overall efficiency decreases because the impression is taken away from a higher-value bidder and transferred to a lower-value bidder. Second, if the impression would have been unsold but is now awarded to a GDN advertiser, overall efficiency increases because the impression is now sold when it otherwise would have gone unsold. In particular, the impact on the overall efficiency of the market is entirely determined by comparing two quantities, on average, when overbidding causes GDN to win an impression it otherwise would not have: (a) the value of the winning advertiser without Project Bernanke (\$0 if the impression would have been unsold), and (b) the highest GDN bid for this impression. If (a) exceeds (b) on average, then the overall market efficiency decreases. If (b) exceeds (a) on average, then the overall market is more efficient. Note, in particular, that because the value of the winning advertiser without Project Bernanke is always higher than the payment AdX would have received without Project Bernanke, it is possible for the overall market efficiency to decrease even if publisher revenue on average increases.

F. Impact of Projects Bernanke and Global Bernanke on Exchanges

262. When combined with (Enhanced) Dynamic Allocation, Projects Bernanke and Global Bernanke enabled AdX to have a higher win rate, which would cause other exchanges to have a lower win rate.³⁹⁸ Under (Enhanced) Dynamic Allocation, AdX won every impression for which it exceeds its reserve. Projects Bernanke and Global Bernanke inflate the highest bid to AdX, making AdX more likely to clear its reserve. Therefore, AdX is more likely to win. This holds both when other exchanges participate in the waterfall process (because fewer impressions will even continue down the waterfall if AdX wins them early), and when other exchanges participate in header bidding (because AdX is now more likely to exceed the highest header bidding bid, which acts as AdX's reserve). Since AdX is more likely to win, this directly implies that other exchanges are more likely to lose, due to the fixed number of impressions.³⁹⁹

263. When AdX participates in a simultaneous auction with other exchanges, Projects Bernanke and Global Bernanke could cause AdX to win more or fewer impressions. When AdX uses a second-price auction to participate in an auction with other exchanges, such as in Exchange Bidding, it matters not only whether AdX clears its reserve, but also at what price it clears. A higher clearing price would cause AdX to win more often, and a lower clearing price would cause AdX to win less often. Under Projects Bernanke and Global Bernanke, the overbidding aspect causes AdX to have a higher clearing price, and therefore would cause AdX to win more often. On the other hand, the collusion aspect causes AdX to have a lower clearing price, and therefore would cause AdX to win less often.⁴⁰⁰ Because impacts are possible in both directions, AdX would sometimes have a higher clearing price and sometimes have a lower clearing price.⁴⁰¹

³⁹⁸ This conclusion assumes that publishers did not significantly inflate AdX's price floor due to Projects Bernanke and Global Bernanke. This seems plausible, due to (a) Projects Bernanke and Global Bernanke were never disclosed and (b) for publishers who let (Enhanced) Dynamic Allocation set AdX's price floor without an additional 'boost', AdX's price floor would be the maximum Value/temporary CPM of other exchanges/guaranteed line items, which are not impacted by GDN's bids on Adx.

³⁹⁹ In particular, exchanges that participate via header bidding are more likely to have their bids superseded by AdX, and exchanges that participate via the waterfall are more likely to see impressions taken by AdX before having an opportunity to solicit bids.

⁴⁰⁰ Note that the collusion aspects cannot cause AdX to fail to meet its reserve. But conditioned on AdX meeting its reserve, it could cause AdX to have a lower clearing price, which is exactly how GDN benefits.

⁴⁰¹ Note that under (Enhanced) Dynamic Allocation, the magnitude of AdX's clearing price does not impact whether it wins the impression. It only matters whether AdX beats its reserve or not.

G. Project First-Price Bernanke

264. AdX eventually switched to a first-price auction,⁴⁰² which renders the particular collusion mechanics of old Project Bernanke obsolete. This is because the first-price auction is pay-your-bid, hence dropping GDN's second highest bid does not impact the auction at all. The general framework of colluding and overbidding still apply, but the precise mechanics differ. I provide more details on this in Appendix H.

265. As I previously discussed, first-price auctions are not truthful. In fact, for any bid b less than the bidder's value v , it is always better to bid b instead of v .⁴⁰³ However, an economic approach called the "Revelation Principle"⁴⁰⁴ allows an intermediary to make a first-price auction truthful for participants. Intuitively, it works in the following way: The intermediary first comes up with a device that takes in a bidder's value as an input and calculates the optimal bid. The intermediary then tells the bidders to report their true values and assures them that if they win, they will be charged their minimum bid to win. But the intermediary does not submit the true values of the bidders to the auction on their behalf, and instead submits bids calculated by the optimization device. The auctioneer then executes a first-price auction with those bids. From the perspective of the advertisers, this is a truthful auction, because they always pay their minimum bid to win. But potentially there is a mismatch in payments since the minimum bid to win (what is charged to the winning bidder) might differ from what is calculated as the optimal bid from the highest value submitted by the bidders (what is paid to the auctioneer). If the device is excellent at bid optimization, these will perfectly balance out on average. If not, there can be a benefit or loss to either party. For the rest of this analysis, I assume that the bid optimizer is excellent.⁴⁰⁵

266. First-Price Project Bernanke has three components: (a) a bid optimizer for GDN users that makes their participation in AdX's first-price auction truthful, (b) collusion among GDN bidders, which increases GDN's payoff at the expense of publishers' revenue, (c) overbidding, which lowers GDN's and increases publishers' revenue. In comparison to Projects Bernanke and Global

⁴⁰² See Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. <https://web.archive.org/web/20240122142404/https://blog.google/products/admanager/update-first-price-auctions-google-ad-manager/>

⁴⁰³ To see this, observe that bidding v guarantees a payoff of 0, no matter what since it either leads to losing the item, or winning the item and paying the bid. Bidding b that is less than v instead guarantees a payoff no worse than 0 since it either leads to losing the item or winning the item and paying less than the value.

⁴⁰⁴ See Roger B. Myerson. "Incentive Compatibility and the Bargaining Problem." *Econometrica* vol. 47, no. 1. 1979. pg. 61–73.

⁴⁰⁵ Shortcomings of the bid-optimizer are certainly relevant for thinking through the impacts of First-Price Project Bernanke, but it is not relevant to the conclusions I draw based on collusion and overbidding alone.

Bernanke, (b) and (c) are conceptually similar but implemented via different mechanics (due to the different mechanics between first- and second-price auctions).

267. First-Price Project Bernanke carries the same motivation as Project Bernanke. GDN bidders could collude in AdX's auction to increase GDN's revenue at the expense of the publisher, but this hurts publishers' revenues. First-Price Project Bernanke again observes that overbidding has the opposite effect of lowering GDN's payoff but helping publishers' revenues, although it causes collateral damage to non-GDN advertisers. Additionally, there is also an added complication due to intermediating AdX's first-price auction to make it truthful.

268. Under First-Price Project Bernanke, GDN manipulated advertisers' bids before sending them to AdX in the following manner:^{406, 407, 408}

a. [REDACTED]

⁴⁰⁶ First Price Project Bernanke was launched in 2019. See GOOG-AT-MDL-008842383 at -88. August 5, 2023. "Declaration of Nirmal Jayaram." ("Google updated the Bernanke algorithms in 2019 to be compatible with the Unified First Price Auction. The updated version of Bernanke was sometimes referred to within Google as 'Alchemist.'")

⁴⁰⁷ GOOG-DOJ-AT-02224828. March 2019. "The Alchemist." [REDACTED]

[REDACTED] My best guess is that the core ideas have not changed.

⁴⁰⁸ This was also called "The Alchemist." [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

269. To illustrate how First-Price Project Bernanke works, imagine that there are ten total bidders, and each is believed to have a valuation uniformly in the range [0,10]. If all players use an ad buying tool with a 15% take rate, the best course of action is for each ad buying tool to submit a bid of $0.9 * 0.85 * v$ to AdX when their advertiser's value is v .⁴¹² If an ad buying tool wanted to modify the auction process to make it truthful without also implementing First-Price Project Bernanke, they could tell their advertisers that when an advertiser submits a bid of b to the tool, the tool submits a bid of $0.9 * 0.85 * b$ on their behalf to AdX. If the tool wins, and the highest other bid is H , then the tool pays $0.9 * 0.85 * b$ to AdX, and charges the winning advertiser $H / (0.9 * 0.85)$ which corresponds to the minimum bid the advertiser could have submitted to their ad buying tool and still had the winning bid. While a single auction is unlikely to result in a take rate of exactly 15%, this will average out to a take rate of 15% over time.

270. First-Price Project Bernanke additionally implements collusion among its advertisers. Say for example that nine of the advertisers use GDN. [REDACTED]

⁴¹² This follows as it is a Bayes-Nash equilibrium for each advertiser to bid $0.9 * v$, and their behavior is not impacted by what take rate the ad buying tools use, provided they all use the same take rate.

- [REDACTED]
- [REDACTED]
- a. Collusive aspect. In the above example, without any collusion, each ad buying tool would optimally submit a bid of $0.9*0.85*v$ on behalf of an advertiser with value v . This is because that advertiser has *uncertainty surrounding the other values of all nine other bidders* and can't shade their bid by too much and still hope to win. If instead, nine of the advertisers get together through GDN and determine who is the highest bidder, now there is only *uncertainty surrounding the other value of just one other bidder*, and it is safe to bid-shade more aggressively. In both a first- and second-price auction, the first step to successful collusion among a team of bidders is to figure out who is the highest bidder. In a second-price auction, the remaining bidders simply drop out in order to lower the winner's eventual payment, but the winning bidder need not adjust their bid as the second-price auction is truthful. In a first-price auction, dropping out alone does not help the winning bidder (because they will still pay their bid if they win, and they will still win if and only if their bid exceeds the highest non-colluding bid). But now that the highest colluding bidder *knows that they face less competition*, they can shade their bid more aggressively to get higher payoff. In the example, this manifests with an optimal bid-shade of 50% versus 90%.⁴¹³
 - b. Overbidding aspect. In the above example, overbidding represents that the ad buying tool submits the optimal bid for a value of $\alpha*b$, instead of the optimal bid for b . In this example, to implement overbidding without collusion, the ad buying tool would submit a bid of $0.9*0.85*\alpha*b$. With collusion and overbidding, the ad buying tool submits a bid of $0.5*0.85*\alpha*b$. This aspect is quite similar to overbidding in a second-price auction, after accounting for the role of the ad buying tool in implementing the Revelation Principle.

⁴¹³ Here is a simpler, but less realistic example. Imagine instead that all ten bidders collude in a first-price auction. Once they determine the highest bidder amongst themselves, that bidder should optimally submit a bid of just a penny, and the others will drop out. Note that if the highest bidder submits their original optimal bid of $0.9*0.85*v$, then they do not benefit simply because the colluders drop out – the highest bidder must further use this information to re-optimize their bid (which in this case is just a penny, as there is no remaining competition). This reasoning qualitatively extends to the chosen example, and indeed all instances, although the required mathematics is more involved.

271. When AdX participates in a simultaneous auction with other exchanges, First-Price Project Bernanke could cause AdX to win more or fewer impressions. When AdX participates in an auction with other exchanges, such as in Exchange Bidding, it matters not only whether AdX clears its reserve, but also at what price it clears. A higher clearing price would cause AdX to win more often, and a lower clearing price would cause AdX to win less often. In a first-price auction, the clearing price is the highest bid, and therefore AdX's clearing price would increase or decrease based on whether Project First-Price Bernanke causes GDN's highest submitted bid to increase or decrease. The overbidding aspect causes GDN to submit a higher bid. On the other hand, the collusion aspect causes GDN to submit a lower bid.⁴¹⁴ Because impacts are possible in both directions, AdX would sometimes have a higher clearing price and sometimes have a lower clearing price.

IX. CONDUCT ANALYSIS: RESERVE PRICE OPTIMIZATION

272. In this section, I provide an analysis of Google's Reserve Price Optimization (RPO)⁴¹⁵ conduct. I demonstrate that RPO leads to higher revenue for Google's ad exchange AdX and explain the mechanisms through which it leads to lower payoff to advertisers and could lead to lower revenue for some publishers. Furthermore, I outline how it impacts publisher and advertiser behavior. The negative effects of RPO to advertiser payoff, and possibly some publishers' revenues, at least partially stem from Google's concealment of the conduct.

273. In particular, it is my opinion that Google concealed information that is material to both publishers and advertisers during the period RPO was concealed. It is also my opinion that even after RPO was revealed, publishers might set suboptimal reserves on any impression for which RPO is a possibility.⁴¹⁶

274. Under RPO, AdX used data available to them (prior to seeing live bids) to calculate per-buyer reserve prices⁴¹⁷ that it believed would optimize AdX's revenue.⁴¹⁸ AdX then used these

⁴¹⁴ Note that the collusion aspect cannot cause GDN to fail to submit a bid above AdX's reserve. But conditioned on meeting AdX's reserve, GDN could submit either a higher or lower bid.

⁴¹⁵ There seems to be other programs that are called RPO previously. However, they substantially differ from the conduct I am discussing here. The main difference between those programs and this conduct is that this conduct sets per-buyer reserve prices.

⁴¹⁶ That is, unless a publisher knows whether RPO activates on a particular impression, and if so what reserve RPO would set, I would expect publishers to lack sufficient information to set a profit-maximizing reserve.

⁴¹⁷ GOOG-NE-13204729 at -36. August 17, 2015. "AdX Dynamic Price." ("Buyer Specific Reserve Prices. Different buyers may get different reserve prices.")

⁴¹⁸ GOOG-NE-06151351 at -52. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers." ("[RPO] generate[s] a histogram of historical bids and transaction prices [...] pick[s] a reserve price that maximizes predicted revenue.")

reserves in its own auction instead of the reserves set by the publisher, although this reserve was always at least as large as the reserve set by the publisher.⁴¹⁹ The program was launched in phases between April and October 2015.⁴²⁰ Initially, Google did not announce this program to its customers.⁴²¹ Later, Google announced the program to its customers under the name “optimized pricing” on May 12th, 2016, more a year after its initial rollout.⁴²² Publishers were not allowed to opt out of the program.⁴²³ The program was deprecated in 2019 with the switch of AdX to the first-price auction format.⁴²⁴

275. Internal Google documents suggests that RPO relies on an algorithmic optimization that “set[s] optimized reserve prices in AdX auction[s]” to “increase the revenue for publishers” via “model[ing] effect of various reserve prices” and then “pick[ing] the best one.”⁴²⁵ Importantly, this tool aims to set the reserve price just below what the highest bidder is willing to pay,⁴²⁶ by coming up with an empirical estimate of this willingness to pay, which was assumed to be equal their bid due to the truthfulness of the second-price auction.⁴²⁷ An internal Google document states that the goal of RPO was to “select a reserve price as close to the anticipated first price as possible in order to trade buyer for seller surplus.”⁴²⁸ If AdX has sufficient data to form an accurate prediction of the maximum advertiser value v , the optimal reserve price to set is exactly v .

Different Google internal documents outline different strategies AdX used to employ the data they have to best estimate the RPO reserve prices. Which data was used and how data was processed are not relevant to the conclusions I provide below.

⁴¹⁹ GOOG-NE-03640022 at -2. “AdX Managed Reserves.” (“Currently RPO can only raise reserve prices.”) Notice that optimizing publisher revenue and AdX revenue are equivalent since AdX revenue corresponds to 20% of the publisher revenue.

⁴²⁰ See, e.g., GOOG-NE-06151351 at -52. November 12, 2015. Email thread, “Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers.” (“Between April and October we launched and improved new systems to dynamically set auction reserve prices for AdX sellers.”)

⁴²¹ GOOG-NE-09485306 at -432. December 18, 2017. “OLD – New Ad Manager Indirect Notes.” (“We are not commercializing this externally for now.”)

⁴²² See Jonathan Bellack. “Smarter optimizations to support a healthier programmatic market” (May 12, 2016). Accessed on May 31, 2024.

<https://web.archive.org/web/20200929015943/https://blog.google/products/admanager/smarter-optimizations-to-support/>

⁴²³ GOOG-NE-06842715 at -18. May 10, 2016. “AdX Auction Optimizations.” (“No opt-out possible.”)

⁴²⁴ GOOG-AT-MDL-000987708 at -8. April 9, 2021. “PM Perspective on 1P RPO.” (“When we transitioned to a 1st price auction and launched unified pricing rules in September 2019, we had to turn off 2P RPO since it was designed to work in a 2nd price auction (duh.”)

⁴²⁵ GOOG-NE-13204729 at -30. August 17, 2015. “AdX Dynamic Price.”

⁴²⁶ When there is sufficient data to predict the highest bidder’s willingness to pay exactly, the goal is indeed to set the reserve price just below this. Often there is insufficient data to predict the highest bidder’s willingness to pay exactly, and in these cases Reserve Price Optimization instead aims to set the optimal reserve given the information it has.

⁴²⁷ GOOG-NE-13204729 at -34. August 17, 2015. “AdX Dynamic Price.” (the slide titled “How to Guess the Top Bid” explains how Google thought about estimating the highest bid.)

⁴²⁸ GOOG-NE-03640022 at -2. “AdX Managed Reserves.”

276. To illustrate how RPO works, imagine an impression arrives for a user who is over the age of 25, likes toys, and lives in Plano, TX. This coarse targeting data is sufficient for the publisher to estimate that such users on average fetch \$5, although certainly there is high variance depending on more fine-grained third-party cookie data. AdX is called on this impression and learns from third-party cookie data that this user further has no children and is working as an unpaid intern. Under RPO, AdX would guess from the third-party cookie data that this impression would not sell for more than \$1, so it leaves the publisher's reserve of \$5 intact. When AdX solicits bids, the highest bid is \$2, and the impression does not clear. Another impression arrives for a user who is over the age of 25, likes toys, and lives in Plano, TX. Now imagine the same setting, but the third-party cookie data shows this user has two children, buys toys for them frequently, and lives in a wealthy neighborhood. Under RPO, AdX would guess from the third-party cookie data that this impression would fetch \$50, so it increases the reserve to \$50. In the AdX auction, perhaps the two highest bids are \$60 and \$55. In this case, with or without RPO, AdX's auction clears at \$55, netting \$11 for AdX and \$44 for the publisher. Perhaps instead the two highest bids are \$60 and \$45. In this case, with RPO, AdX's auction clears at \$50, netting \$10 for AdX and \$40 for the publisher. Without RPO, AdX's auction would have cleared at \$45, netting \$9 for AdX and \$36 for the publisher. In this case, RPO increases AdX's and publisher's revenue, but the advertiser suffers from a higher payment. Note that in all cases, the advertiser suffers under RPO due to facing a higher reserve. In these particular examples, both publisher and AdX profit modestly under RPO, although it is possible for both to profit significantly (at the expense of the advertiser's payoff, in case that the advertiser might have paid way below their value without RPO), or to have significantly decreased profit (at significant cost to the advertiser as well, if RPO accidentally sets too high of a reserve and the impression no longer clears).

A. Impact of Reserve Price Optimization on Publishers

277. Under general circumstances, because RPO can only increase the publisher reserve, it could increase publishers' payoff by (a) generating greater expected revenue from sales if the publisher believes that Google is better at optimizing revenues than themselves, perhaps due to better data and more sophisticated algorithms and (b) causing AdX to be less likely to win impressions, which might improve the publisher's payoff if AdX ads are typically of low quality.⁴²⁹

⁴²⁹ A higher reserve for an auction leads to a lower probability of the item being sold in either a first- or second-price auction.

278. Google internal documents reveal that RPO increased publisher revenues. A launch email states that the program generated an annual revenue increase [REDACTED] for publishers.⁴³⁰ Increased revenues notwithstanding, it is still my opinion that concealing RPO from publishers conceals material information from publishers, for the reasons outlined below.

1) RPO can prevent publishers from optimizing their revenue

279. With RPO, Google concealed material information from publishers. As noted in Section II, reserve prices are material to a publisher's revenue. As one example, if Google is good at optimizing reserves via RPO, a publisher may wish to lower the reserve it sets on AdX in order to give AdX greater flexibility in optimizing its reserve, which would lead to greater revenues for both AdX and the publisher. But, via concealing RPO, Google prevented the publishers from effectively optimizing revenue. In fact, even after RPO is disclosed, publishers would still face challenges setting optimal reserves under RPO. For example, if a publisher wishes to lower the reserve it sets on AdX to give RPO more flexibility, they would want to know exactly on which auctions RPO is active. If a publisher prefers to trust their own optimization over Google's, they would further want to know not only whether RPO is active, but also exactly what reserve RPO is setting (so they can set what they consider to be the optimal reserve exceeding RPO's).

280. However, there are also reasons why some publishers might prefer outcomes without RPO than with RPO. For example, a publisher whose goal is to maximize the number of sold ads (as opposed to maximizing their revenue) would prefer outcomes without RPO as opposed to with RPO. Such publishers would want a low reserve to achieve their goal of optimizing sold inventory. As another example, a publisher may have a data team they have invested in that sets reserve prices based on tailor-made algorithms for their user and advertiser base. As a result, this publisher may believe that their reserve generates greater revenues than what would result from RPO and prefer the outcomes with their own algorithm instead of RPO. As another possible example, a publisher might have a small data team optimizing the AdX reserve under the assumption that this is the reserve set in AdX, and that when an impression fails to sell through AdX it is because no bid above the reserve was received from AdX. This would be an incorrect assumption under RPO, because it could instead be that such a bid was received but filtered by AdX for not clearing the higher reserve set by RPO. Hence, the publisher data team might have

⁴³⁰ [REDACTED]

[REDACTED]

[REDACTED]

corrupted data, and as a result fail to optimize their reserve on AdX in the future based on this data, or mistakenly use this corrupted data to poorly set reserves for other exchanges.⁴³¹

B. Impact of Reserve Price Optimization on Advertisers

281. RPO would lead to a payoff loss for advertisers since it leads to both a decrease in impressions won and an increase in the average price paid for impressions won. This is because the advertisers face higher reserves, which can lead to:

- a. A decrease in the advertiser win rate, since RPO increases publisher reserves, and a higher reserve leads to a lower probability of the impression being cleared in a first- or second-price auction. As a result, advertisers might lose impressions that they would have won otherwise.
- b. An increase in the average clearing price, during periods when AdX ran a second-price auction, since a higher reserve leads to a higher clearing price when the second highest bid is the reserve itself. Other auctions where there are two bids above the reserve either remain unaffected (if the increased reserve is still below the second highest bid) or RPO leads to a higher clearing price in these as well (if the increased reserve surpasses the second highest bid). As a result, the publishers must pay more for impressions they win compared to what they would have without RPO.
- c. During periods when both RPO and Exchange Bidding were active, an increased clearing price would cause AdX to submit a higher clearing price to Exchange Bidding, which could cause AdX to win additional impressions over competing exchanges' bids. Note that this reasoning does not apply prior to Exchange Bidding. Prior to Exchange Bidding, AdX won every auction for which it found an advertiser above its reserve due to (Enhanced) Dynamic Allocation.

282. Google internal documents demonstrate this as well. A launch email states that "setting optimized prices on behalf of publishers makes queries more expensive for buyers."⁴³²

⁴³¹ These are all reasons why a publisher could reasonably prefer outcomes without RPO to outcomes with RPO and highlight reasons why the existence of RPO is material to publishers. These reasons also extend to any auction for which a publisher believes RPO is possible, even if RPO is ultimately inactive on that auction.

⁴³² GOOG-NE-06151351 at -53. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers."

- 1) Advertisers would change their bidding behavior had Google revealed RPO during its initial implementation

283. Google concealed vital information from advertisers by concealing RPO. In particular, if Google advertised AdX as a truthful second-price auction,⁴³³ this would encourage advertisers to bid their true values for impressions. However, if RPO was also using past AdX bid data to set reserves for future AdX auctions, the long-term process is no longer truthful. That is, for any particular AdX auction, advertisers indeed maximize their payoff in this individual auction by reporting their true value. However, submitting a high bid equal to an advertiser's true value would cause later AdX reserves to increase, decreasing that advertiser's future payoff in later AdX auctions.⁴³⁴

284. To see why that is the case, imagine an advertiser who has just opened the first Lego shop in Plano, TX is planning to advertiser for their business. Prior to this, the demand for an impression to a user who is over the age of 25, lives in Plano, and likes Legos might be just \$1, because the value to most advertisers is simply just because the user is over the age of 25 and lives in Plano. RPO would learn this over time and set a reserve of \$1 for such users. This advertiser, on the other hand, has the perfect business for this user, and has a value of \$10. The advertiser believes that they are participating in a regular second-price auction with reserve \$1, and therefore bids \$10 for all such impressions, since bidding their value is the best strategy for them in a second-price auction. Initially, they win them all and pay \$1 each, because they are indeed participating in a second-price auction with reserve \$1. Google's RPO algorithm notices, however, that impressions for users over the age of 25, living in Plano, and who like Legos tend to elicit bids closer to \$10, which is much more than the current reserve of \$1. RPO does its job and raises the reserve, eventually getting quite close to \$10. This means that while the advertiser enjoyed a good profit initially, their profit for these impressions is going to converge to \$0 in the long run due to RPO, since they will be paying their value for each impression.

⁴³³ Google claimed that AdX ran a second-price auction up until it switched to the first-price auction format. See Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. <https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/>

⁴³⁴ GOOG-NE-13207530 at -30. August 25, 2015. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

285. Note that each individual auction in isolation is still a second-price auction with reserve, and so an advertiser cannot improve their gain in that auction by shading their bid. However, if the advertiser were aware of RPO, they would have shaded their bid from the beginning. For example, if the advertiser instead consistently submitted a bid of \$2, they would still win every impression, but they would instead have a long run per-impression profit of \$8 since they value the impressions at \$10 but pay \$2.

C. Impact of Reserve Price Optimization on AdX

286. The impact of RPO on AdX is indeterminate and would depend on how effective RPO is. The program can lead to:

- a. A decrease in the AdX win rate, since RPO increases publisher reserves and a higher reserve leads to a lower probability of the impression being cleared in a second-price auction. A lower win rate has a negative effect on AdX revenue.
- b. An increase in the average AdX clearing price in the auctions they clear, since a higher reserve leads to a higher clearing price when the second highest bid is the reserve itself. Other auctions where there are two bids above the reserve either remain unaffected (if the increased reserve is still below the second highest bid) or RPO leads to a higher clearing price in these as well (if the increased reserve surpasses the second highest bid). As a result, average AdX clearing price increases, which has a positive effect on AdX revenue since AdX takes a fee based on the clearing price.
- c. During periods when both RPO and Exchange Bidding were active, an increased clearing price would cause AdX to submit a higher clearing price to Exchange Bidding, which could cause AdX to win additional impressions over competing exchanges' bids. Note that this reasoning does not apply prior to Exchange Bidding. Prior to Exchange Bidding, AdX won every auction for which it found an advertiser above its reserve due to (Enhanced) Dynamic Allocation.

287. RPO is effective when the gains from b (and c, during Exchange Bidding) outweigh the losses from a, and it is reasonable to expect a sophisticated seller with ample data, such as Google, to accomplish this.

288. Google documents suggest that RPO indeed improved Google's revenue. For example, a post-launch email notes that RPO led to an increase in annual revenue [REDACTED]

[REDACTED] for publishers,⁴³⁵ and a March 2016 brief estimates the annual increase in Google revenue due to RPO [REDACTED].⁴³⁶

289. RPO allows AdX to increase the reserve price beyond what the publisher sets. Dynamic Revenue Sharing⁴³⁷ allows AdX to functionally lower the effective reserve price⁴³⁸ below what the publisher sets.^{439, 440} Both conduct together allow AdX the flexibility to adjust the publisher-set reserve in either direction.⁴⁴¹

⁴³⁷ See Section VII for a discussion of Dynamic Revenue Sharing.

⁴³⁸ That is, the reserve price plus the ad exchange take rate.

⁴³⁹ GOOG-NE-13205325 at -36. (comment [5] suggests further that DRSv2 could also lower the take-rate applied to an RPO-set reserve, and not just the publisher-set reserve.)

⁴⁴⁰ In principle, DRS, if active on RPO-set reserves, can also help AdX clear impressions where RPO mistakenly sets a higher reserve price than what is needed. In such a case, DRS can help AdX clear the impression, as long as the amount that RPO is mistaken by is lower than what DRS can achieve by decreasing the AdX take rate to 0%.

⁴⁴¹ The magnitude with which AdX can increase the publisher's reserve using RPO is unbounded, while the magnitude with which AdX can decrease the publisher's reserve using DRS is bounded. AdX will always collect at least the publisher's price floor with a 0% take rate from the winning advertiser.